



Securing applications and their environment

Note

Before using this information, be sure to read the general information under “Notices” on page 1337.

Compilation date: September 24, 2008

© Copyright International Business Machines Corporation 2008.

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

Contents

How to send your comments	xiii
Changes to serve you more quickly	xv
Chapter 1. Overview and new features for securing applications and their environment	1
Security planning overview	1
Chapter 2. Task overview: Securing resources	13
Chapter 3. Setting up and enabling security	15
Migrating, coexisting, and interoperating – Security considerations	15
Interoperating with previous product versions	16
Interoperating with a C++ common object request broker architecture client	18
Migrating custom user registries	20
Migrating trust association interceptors	22
Migrating Common Object Request Broker Architecture programmatic login to Java Authentication and Authorization Service (CORBA and JAAS)	25
Migrating from the CustomLoginServlet class to servlet filters	28
Migrating Java 2 security policy	29
Migrating with Tivoli Access Manager for authentication enabled on a single node	31
Migrating with Tivoli Access Manager for authentication enabled on multiple nodes	32
Preparing for security at installation time	34
Securing your environment before installation	35
Securing your environment after installation	35
Enabling security	37
Administrative security	40
Considering security specific to a multi-node or process Network Deployment environment	42
Application security	43
Java 2 security	44
Enabling security for the realm	53
Testing security after enabling it	69
The Security Configuration Wizard	70
Securing specific application servers	70
Server-level security settings	72
Chapter 4. Configuring multiple security domains	75
Multiple security domains	78
Creating new multiple security domains	93
Deleting multiple security domains	95
Copying multiple security domains	96
Configuring inbound trusted realms for multiple security domains	99
Configure security domains	99
Name	100
Description	100
Assigned Scopes	100
Application Security:	101
Enable application security	101
Java 2 security:	101
Use global security settings	101
Customize for this domain	102
Use Java 2 security to restrict application access to local resources	102
Warn if applications are granted custom permissions	102
Restrict access to resource authentication data	102

User Realm:	102
Trust Association:	103
Interceptors	103
Enable trust association	103
SPNEGO Web Authentication:	103
RMI/IIOP Security:	104
CSlv2 inbound communications	104
CSlv2 outbound communications	104
JAAS Application logins	104
Use global and domain-specific logins	105
JAAS System Logins:	105
System Logins	105
JAAS J2C Authentication Data:	105
Use global and domain-specific entries	105
Authentication Mechanism Attributes:	105
Authorization Provider:	105
Custom properties	105
Web Services Bindings	106
External realm name	106
External realm name	106
Trust all realms	106
Trust all realms (including those external to this cell)	106
Trust realms as indicated below	106
Add External Realm...	106
Security domains collection	107
Maximum rows	107
Retain filter criteria	107
Copy selected domain	107
Copy global security	107
Authentication cache settings	107
Enable authentication cache	107
Cache timeout:	108
Initial cache size:	108
Maximum cache size	108
Use basic authentication cache keys (password one-way hashed):	108
Use custom cache keys:	109
Chapter 5. Authenticating users	111
Selecting a registry or repository	111
Standalone custom registries	114
Configuring local operating system registries	115
Configuring Lightweight Directory Access Protocol user registries	119
Configuring standalone custom registries	142
Managing the realm in a federated repository configuration	170
Local operating system registries	238
Standalone Lightweight Directory Access Protocol registries	242
Federated repositories	246
Selecting an authentication mechanism	247
Lightweight Third Party Authentication	249
Configuring the Lightweight Third Party Authentication mechanism	250
Kerberos (KRB5) authentication mechanism support for security	259
Configuring Kerberos as the authentication mechanism using the administrative console	266
Configuring a Java client for Kerberos authentication	280
RSA token authentication mechanism	282
Configuring the RSA token authentication mechanism	284
Message layer authentication	288

Trust associations	289
Integrating third-party HTTP reverse proxy servers	293
Trust association settings	293
Trust association interceptor collection	294
Trust association interceptor settings	294
Single sign-on	294
Single sign-on using LTPA cookies.	295
Global single sign-on principal mapping	296
Implementing single sign-on to minimize Web user authentications	297
Creating a single sign-on for HTTP requests using SPNEGO Web authentication	300
Creating a single sign-on for HTTP requests using the SPNEGO TAI (deprecated)	318
Configuring single sign-on capability with Tivoli Access Manager or WebSEAL	354
Configuring administrative authentication	369
Job manager security	370
Java Authentication and Authorization Service	371
Java Authentication and Authorization Service authorization	371
Using the Java Authentication and Authorization Service programming model for Web authentication	373
Developing custom login modules for a system login configuration	375
Customizing application login with Java Authentication and Authorization Service	388
Performing identity mapping for authorization across servers in different realms	434
Configuring inbound identity mapping.	435
Configuring outbound mapping to a different target realm	443
Security attribute propagation	448
Default authentication token	452
Propagating security attributes among application servers	453
Using the default authorization token	456
Using the default propagation token	459
Using the default single sign-on token with default or custom token factory	465
Configuring the authentication cache	466
Configuring Common Secure Interoperability Version 2 (CSIV2) inbound and outbound communication settings	466
Configuring Common Secure Interoperability Version 2 inbound communications.	467
Configuring Common Secure Interoperability Version 2 outbound communications	473
Configuring inbound transports	478
Configuring outbound transports	483
Configuring inbound messages	486
Configuring outbound messages	487
Common Secure Interoperability Version 2 and Security Authentication Service (SAS) client configuration	488
Example 1: Configuring basic authentication and identity assertion	493
Example 2: Configuring basic authentication, identity assertion, and client certificates	494
Example 3: Configuring client certificate authentication and RunAs system	496
Example 4: Configuring TCP/IP transport using a virtual private network	497
Authentication protocol for EJB security	499
Authentication protocol support	502
Common Secure Interoperability Version 2 features	503
Identity assertion to the downstream server	503
Identity assertions with trust validation	505
Message layer authentication	505
Chapter 6. Authorizing access to resources	509
Authorization technology	509
Administrative roles and naming service authorization	510
Role-based authorization	516
Administrative roles	518
Authorization providers	522

Delegations	535
Authorizing access to J2EE resources using Tivoli Access Manager	537
Using the built-in authorization provider	538
Enabling an external JACC provider	542
Authorizing access to administrative roles	561
Administrative user roles settings and CORBA naming service user settings	562
Administrative group roles and CORBA naming service groups	563
Assigning users to naming roles	565
Propagating administrative role changes to Tivoli Access Manager	566
migrateEAR utility for Tivoli Access Manager	567
Assigning users from a foreign realm to the admin-authz.xml	570
Fine-grained administrative security	570
New Administrative Authorization Group	573
Administrative Authorization Group collection	574
Creating a fine-grained administrative authorization group using the administrative console	575
Editing a fine-grained administrative authorization group using the administrative console	578
Fine-grained administrative security in heterogeneous and single-server environments	580
Example: Using fine-grained security	581
Chapter 7. Securing communications	589
Secure communications using Secure Sockets Layer	589
Secure Sockets Layer configurations	595
Keystore configurations	604
Dynamic outbound selection of Secure Sockets Layer configurations	607
Central management of Secure Sockets Layer configurations	608
Secure Sockets Layer node, application server, and cluster isolation	609
Default chained certificate configuration	614
Dynamic configuration updates	624
Management scope configurations	625
Certificate management using iKeyman	626
Certificate management	627
Using the retrieveSigners command to enable server to server trust	631
Creating a Secure Sockets Layer configuration	632
SSL certificate and key management	635
SSL configurations for selected scopes	636
SSL configurations collection	637
SSL configuration settings	637
Certificate authority (CA) client configuration	639
Certificate authority (CA) client configuration collections	640
Create chained personal certificate	640
Recovering deleted certificates	641
Renewing a certificate	642
Revoke a CA certificate	642
Using a CA client to create a personal certificate to be used as the default personal certificate	643
Create a CA certificate	644
Developing the WSPKIClient interface for communicating with a certificate authority	645
Creating a custom trust manager configuration	646
Creating a custom key manager	652
Associating a Secure Sockets Layer configuration dynamically with an outbound protocol and remote secure endpoint	658
Quality of protection (QoP) settings	668
ssl.client.props client configuration file	670
Create a CA client	677
Delete a CA client	677
View or modify a CA client	678
Creating a keystore configuration for a preexisting keystore file	679

Configuring a hardware cryptographic keystore	680
Managing keystore configurations remotely	681
Keystores and certificates collection	682
Key store settings	683
Key managers collection	685
Key managers settings	685
Creating a self-signed certificate	686
Replacing an existing self-signed certificate	687
Creating a certificate authority request	688
Certificate request settings	689
Personal certificates collection	689
Self-signed certificates settings	691
Personal certificate requests collection	693
Personal certificate requests settings	694
Extract certificate request	695
Receiving a certificate issued by a certificate authority	696
Replace a certificate	699
Extracting a signer certificate from a personal certificate	700
Extract certificate	701
Extract signer certificate	701
Retrieving signers using the retrieveSigners utility at the client	702
Changing the signer auto-exchange prompt at the client	703
Retrieving signers from a remote SSL port	704
Retrieve from port	705
Adding a signer certificate to a keystore	706
Add signer certificate settings	707
Signer certificates collection	707
Signer certificate settings	708
Adding a signer certificate to the default signers keystore	708
Exchanging signer certificates	710
Keystores and certificates exchange signers	711
Configuring certificate expiration monitoring	711
Manage certificate expiration settings	712
Notifications	714
Notifications settings	715
Key management for cryptographic uses	716
Creating a key set configuration	717
Active key history collection	719
Add key alias reference settings	719
Key sets collection	720
Key sets settings	720
Creating a key set group configuration	722
Example: Retrieving the generated keys from a key set group	723
Example: Developing a key or key pair generation class for automated key generation	724
Key set groups collection	727
Key set groups settings	728
Configuring the Web server plug-in for Secure Sockets Layer	729
Web server plug-in default configuration	730
Chapter 8. Developing extensions to the WebSphere security infrastructure	731
Developing standalone custom registries	731
Example: Standalone custom registries	732
Result.java file	732
UserRegistry.java files	733
Implementing custom password encryption	740
Developing applications that use programmatic security	740

Protecting system resources and APIs (Java 2 security)	741
Developing with programmatic security APIs for Web applications	764
Developing with programmatic APIs for EJB applications	769
Customizing Web application login.	773
Example: Form login	776
Developing servlet filters for form login processing	778
Secure transports with JSSE and JCE programming interfaces	782
Configuring Federal Information Processing Standard Java Secure Socket Extension files	786
Implementing tokens for security attribute propagation	789
Implementing a custom propagation token	789
Implementing a custom authorization token	797
Implementing a custom single sign-on token	806
Implementing a custom authentication token	816
Propagating a custom Java serializable object	825
Developing a custom interceptor for trust associations	828
Trust association interceptor support for Subject creation	832
Enabling a plugpoint for custom password encryption.	834
Plug point for custom password encryption	835
Chapter 9. Auditing the security infrastructure	839
Enabling the security auditing subsystem	840
Security Auditing detail	842
Context object fields	843
Creating security auditing event type filters	846
Auditable security events	847
Event type filter settings	847
Event type filters collection	848
Example: Generic Event Interface	849
Context objects for security auditing	850
Context object fields	850
Configuring security audit subsystem failure notifications	853
Audit monitor collection	854
Audit notification settings	855
Configuring the default audit service providers for security auditing	856
Audit service provider collection.	857
Audit service provider settings	857
Example: Base Generic Emitter Interface	858
Configuring a third party audit service providers for security auditing	858
Example: Base Generic Emitter Interface	859
Configuring audit event factories for security auditing	860
Audit event factory configuration collection.	861
Audit event factory settings	861
Example: Generic Event Factory Interface	862
Protecting your security audit data.	863
Encrypting your security audit records	864
Signing your security audit records	865
Audit encryption keystores and certificates collection	866
Audit record encryption configuration settings.	866
Audit record signing configuration settings	867
Audit record keystore settings	868
Using the audit reader	869
Example: Audit Event Outcome Codes	871
Chapter 10. Configuring security with scripting	875
Enabling and disabling security using scripting	876
Enabling and disabling Java 2 security using scripting	877

Configuring multiple security domains using scripting	879
Configuring security domains using scripting	880
Configuring local operating system user registries using scripting	881
Configuring custom user registries using scripting	883
Configuring JAAS login modules using scripting	885
Configuring Common Secure Interoperability authentication using scripting	886
Configuring trust association using scripting	888
Mapping resources to security domains using scripting	889
Removing resources from security domains using scripting	890
Removing security domains using scripting	890
Removing user registries using scripting	891
SecurityDomainCommands command group for the AdminTask object	892
SecurityConfigurationCommands command group for the AdminTask object	899
SecurityRealmInfoCommands command group for the AdminTask object.	939
NamingAuthzCommands command group for the AdminTask object	945
Utility scripts	950
Configuring the JACC provider for Tivoli Access Manager using the wsadmin utility	955
Securing communications using the wsadmin tool	957
Creating an SSL configuration at the node scope using scripting.	957
Automating SSL configurations using scripting	960
Updating default key store passwords using scripting	963
Configuring certificate authority client objects using the wsadmin tool	963
Administering certificate authority clients using the wsadmin tool.	965
Setting a certificate authority certificate as the default certificate using the wsadmin tool	967
Creating certificate authority (CA) personal certificates using the wsadmin tool	969
Revoking certificate authority personal certificates using the wsadmin tool	971
CAClientCommands command group for the AdminTask object	972
Creating self-signed certificates using scripting	976
keyManagerCommands command group for the AdminTask object	978
KeyStoreCommands command group for the AdminTask object	981
SSLConfigCommands command group for the AdminTask object	990
SSLConfigGroupCommands group for the AdminTask object.	1000
TrustManagerCommands command group for the AdminTask object	1004
KeySetCommands command group for the AdminTask object	1007
KeyReferenceCommands command group for the AdminTask object.	1013
KeySetGroupCommands command group for the AdminTask object	1017
DynamicSSLConfigSelections command group for the AdminTask object	1020
PersonalCertificateCommands command group for the AdminTask object	1022
WSCertExpMonitorCommands command group for the AdminTask object	1038
SignerCertificateCommands command group for the AdminTask object	1043
CertificateRequestCommands command group of the AdminTask object	1049
Enabling authentication in the file transfer service using scripting	1053
Propagating security policy of installed applications to a JACC provider using wsadmin scripting	1054
JACCUtilityCommands command group for the AdminTask object.	1055
Configuring custom adapters for federated repositories using wsadmin	1056
Disabling embedded Tivoli Access Manager client using wsadmin.	1059
Configuring security auditing using scripting	1060
Configuring audit service providers using scripting	1061
Configuring audit event factories using scripting	1062
Configuring auditable events using scripting	1064
Enabling security auditing using scripting	1065
Configuring security audit notifications using scripting	1067
Encrypting security audit data using scripting	1068
Signing security audit data using scripting	1070
AuditKeyStoreCommands command group for the AdminTask object.	1071
AuditEmitterCommands for the AdminTask object.	1077

AuditSigningCommands command group for the AdminTask object	1087
AuditEncryptionCommands command group for the AdminTask object	1093
AuditEventFactoryCommands for the AdminTask object	1109
AuditFilterCommands command group for the AdminTask object	1117
AuditNotificationCommands command group for the AdminTask object	1131
AuditPolicyCommands command group for the AdminTask object	1142
AuditEventFormatterCommands command group for the AdminTask object	1151
AuditReaderCommands command group for the AdminTask object	1152
SSLMigrationCommands command group for the AdminTask object	1155
IdMgrConfig command group for the AdminTask object	1158
IdMgrRepositoryConfig command group for the AdminTask object	1163
IdMgrRealmConfig command group for the AdminTask object	1210
WIMManagementCommands command group for the AdminTask object	1218
DescriptivePropCommands command group for the AdminTask object	1230
ManagementScopeCommands command group for the AdminTask object	1233
AuthorizationGroupCommands command group for the AdminTask object	1235
ChannelFrameworkManagement command group for the AdminTask object	1247
SpnegoTAICommands group for the AdminTask object (deprecated)	1250
The Kerberos configuration file	1255
SPNEGO Web authentication configuration commands.	1258
SPNEGO Web authentication filter commands	1259
Kerberos authentication commands	1262
LTPA_LDAPSecurityOn and LTPA_LDAPSecurityOff command usage	1265
Chapter 11. Configuring denial of service protection for the proxy server	1267
Denial of Service Protection.	1267
Maximum request body buffer chunk size.	1268
Maximum response body buffer chunk size	1268
Chapter 12. Tuning the security properties for the DMZ Secure Proxy Server for IBM	
WebSphere Application Server	1269
DMZ Secure Proxy Server for IBM WebSphere Application Server start up user permissions	1270
DMZ Secure Proxy Server for IBM WebSphere Application Server routing considerations	1270
DMZ Secure Proxy Server for IBM WebSphere Application Server administration options	1271
Error handling security considerations for the DMZ Secure Proxy Server for IBM WebSphere	
Application Server	1271
Proxy security level properties	1272
Current security level	1272
Administration	1272
Routing	1272
Start-up permissions	1273
Custom Error Page Policy	1273
Chapter 13. Tuning, hardening, and maintaining	1275
Tuning security configurations	1275
Secure Sockets Layer performance tips	1278
Tuning security	1280
Hardening security configurations.	1281
Securing passwords in files	1281
Encoding passwords in files.	1282
Enabling custom password encryption	1285
Chapter 14. Troubleshooting security configurations	1289
Security components troubleshooting tips.	1289
Security configuration and enablement errors	1300
Security enablement followed by errors	1303

Access problems after enabling security	1310
Secure Sockets Layer errors	1314
Errors configuring Secure Sockets Layer encrypted access	1318
Single sign-on configuration troubleshooting tips	1320
Security authorization provider troubleshooting tips	1323
SPNEGO trust association interceptor (TAI) troubleshooting tips (deprecated)	1327
Appendix. Directory conventions	1333
Notices	1337
Trademarks and service marks	1339

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 securing applications and their environment

Use the links provided in this topic to learn more about the security infrastructure.

What is new for security specialists

This topic provides an overview of new and changed features in security.

Security

This topic describes how IBM® WebSphere® Application Server provides security infrastructure and mechanisms to protect sensitive Java™ Platform, Enterprise Edition (Java EE) resources and administrative resources and to address enterprise end-to-end security requirements on authentication, resource access control, data integrity, confidentiality, privacy, and secure interoperability.

“Security planning overview”

Several communication links are provided from a browser on the Internet, through Web servers and product servers, to the enterprise data at the back-end. This topic examines some typical configurations and common security practices. WebSphere Application Server security is built on a layered security architecture. This section also examines the security protection offered by each security layer and common security practice for good quality of protection in end-to-end security.

Samples

The Samples Gallery offers:

- **Login - Form Login**

The Form Login Sample demonstrates a very simple example of how to use the login facilities for WebSphere Application Server to implement and configure login applications. The Sample uses the Java Platform, Enterprise Edition (Java EE) form-based login technology to customize the look and feel of the login screens. It uses servlet filters to log the user information and the date information. The Sample finishes the session by using the form-based logout function, an IBM extension to the Java EE specification.


- **Login - JAAS Login**

The JAAS Login Sample demonstrates how to use the Java Authentication and Authorization Service (JAAS) with WebSphere Application Server. The Sample uses server-side login with JAAS to authenticate a real user to the WebSphere security run time. Based upon a successful login, the WebSphere security run time uses the authenticated Subject to perform authorization checks on a protected stateless session enterprise bean. If the Sample runs successfully, it displays all the principals and public credentials of the authenticated user.

Security planning overview

When you access information on the Internet, you connect through Web servers and product servers to the enterprise data at the back end. This section examines some typical configurations and common security practices.

This section also examines the security protection that is offered by each security layer and common security practice for good quality of protection in end-to-end security. The following figure illustrates the building blocks that comprise the operating environment for security within WebSphere Application Server:



The following information describes each of the components of WebSphere Application Server security, Java security, and Platform security that are illustrated in the previous figure.

WebSphere Application Server security

WebSphere security

WebSphere Application Server security enforces security policies and services in a unified manner on access to Web resources, enterprise beans, and JMX administrative resources. It consists of WebSphere Application Server security technologies and features to support the needs of a secure enterprise environment.

Java security

Java Platform, Enterprise Edition (Java EE) security application programming interface (API)

The security collaborator enforces Java Platform, Enterprise Edition (Java EE)-based security policies and supports Java EE security APIs.

CORBA security (CSlv2)

Any calls made among secure Object Request Brokers (ORB) are invoked over the Common Security Interoperability Version 2 (CSlv2) security protocol that sets up the security context and the necessary quality of protection. After the session is established, the call is passed up to the enterprise bean layer.

Java 2 security

The Java 2 Security model offers fine-grained access control to system resources including file system, system property, socket connection, threading, class loading, and so on. Application code must explicitly grant the required permission to access a protected resource.

Java Virtual Machine (JVM) 5.0

The JVM security model provides a layer of security above the operating system layer. For example, JVM security protects the memory from unrestricted access, creates exceptions when errors occur within a thread, and defines array types.

Platform security

Operating system security

The security infrastructure of the underlying operating system provides certain security services for WebSphere Application Server. These services include the file system security support that secures sensitive files in the product installation for WebSphere Application Server. The system administrator can configure the product to obtain authentication information directly from the operating system user registry.

The security infrastructure of the underlying operating system provides certain security services for WebSphere Application Server. The operating system identity of the servant controller, and daemon Started Task, as established by the STARTED profile, is the identity that is used to control access to system resources such as files or sockets. Optionally, the operating system security can provide authentication services using the User Registry of local operating system, and/or authorization services using SAF Authorization for the WebSphere Administration console and for applications running under the application server.

In addition to knowledge of Secure Sockets Layer (SSL) and Transport Layer Security (TLS), the administrator must be familiar with System Authorization Facility (SAF) and Resource Access Control Facility (RACF®), or an equivalent SAF based product.

The identity and verification of users can be managed by using a Local Operating System as the User Registry, RACF or equivalent SAF base product. Alternatively, an LDAP, Custom, or Federated User Registry can be used.

WebSphere can be configured to use SAF Authorization, which will use RACF or an equivalent SAF based product to manage and protect users and group resources. Alternatively, WebSphere can be configured to use WebSphere Authorization or a JACC External Authorization Provider.

When using either Local Operating System as the User Registry and/or using SAF Authorization, security auditing is an inherit feature of RACF or the equivalent SAF based products.

Network security

The Network Security layers provide transport level authentication and message integrity and confidentiality. You can configure the communication between separate application servers to use Secure Sockets Layer (SSL). Additionally, you can use IP Security and Virtual Private Network (VPN) for added message protection.

WebSphere Application Server Network Deployment installation

Note: A node agent instance exists on every computer node.

Each product application server consists of a Web container, an Enterprise Java Beans (EJB) container, and the administrative subsystem.


The WebSphere Application Server deployment manager contains only WebSphere Application Server administrative code and the administrative console.

The administrative console is a special Java EE Web application that provides the interface for performing administrative functions. WebSphere Application Server configuration data is stored in XML descriptor files, which must be protected by operating system security. Passwords and other sensitive configuration data can be modified using the administrative console. However, you must protect these passwords and sensitive data. For more information, see “Encoding passwords in files” on page 1282.

The administrative console Web application has a setup data constraint that requires access to the administrative console servlets and JavaServer Pages (JSP) files only through an SSL connection when administrative security is enabled.

In WebSphere Application Server Version 6.0.x and earlier, the administrator console HTTPS port was configured to use `DummyServerKeyFile.jks` and `DummyServerTrustFile.jks` with the default self-signed certificate. The dummy certificates and keys must be replaced immediately after WebSphere Application Server installation; the keys are common in all of the installation and are therefore insecure. WebSphere Application Server Version 6.1 provides integrated certificate and key management, which generate distinct private key and self-signed certificate with embedded server host name to enable host name verification. WebSphere Application Server Version 6.1 also enables integration with external certificate (CA) authority to use CA-issued certificates. The WebSphere Application Servers Version 6.1 installation process provides an option to enable administrative security during installation. As a result, a WebSphere Application Server process is secured immediately after installation. WebSphere Application Server Version 7.0 extends the embedded certificate management capabilities by creating a chained certificate (personal certificate signed by a root certificate) to enable refresh of the personal certificate without affecting the trust established. It also enables tailoring of the certificate during profile creation (you can import your own or change the distinguished name (DN) of the one created by default) as well as the ability to change the default keystore password.

The following figure shows a typical multiple-tier business computing environment.



Administrative security

WebSphere Application Servers interact with each other through CSv2 and Secure Authentication Services (SAS) security protocols as well as the HTTP and HTTPS protocols.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

You can configure these protocols to use Secure Sockets Layer (SSL) when you enable WebSphere Application Server administrative security. The WebSphere Application Server administrative subsystem in every server uses SOAP, Java Management Extensions (JMX) connectors and Remote Method Invocation

over the Internet Inter-ORB Protocol (RMI/IIOP) JMX connectors to pass administrative commands and configuration data. When administrative security is disabled, the SOAP JMX connector uses the HTTP protocol and the RMI/IIOP connector uses the TCP/IP protocol. When administrative security is enabled, the SOAP JMX connector always uses the HTTPS protocol. When administrative security is enabled, you can configure the RMI/IIOP JMX connector to either use SSL or to use TCP/IP. It is recommended that you enable administrative security and enable SSL to protect the sensitive configuration data.

When administrative security is enabled, you can disable application security at each individual application server by clearing the **Enable administrative security** option at the server level. For more information, see “Securing specific application servers” on page 70. Disabling application server security does not affect the administrative subsystem in that application server, which is controlled by the security configuration only. Both administrative subsystem and application code in an application server share the optional per server security protocol configuration.

Security for Java EE resources

Security for Java EE resources is provided by the Web container and the EJB container. Each container provides two kinds of security: declarative security and programmatic security.

In declarative security, an application security structure includes network message integrity and confidentiality, authentication requirements, security roles, and access control. Access control is expressed in a form that is external to the application. In particular, the deployment descriptor is the primary vehicle for declarative security in the Java EE platform. WebSphere Application Server maintains Java EE security policy, including information that is derived from the deployment descriptor and specified by deployers and administrators in a set of XML descriptor files. At runtime, the container uses the security policy that is defined in the XML descriptor files to enforce data constraints and access control.

When declarative security alone is not sufficient to express the security model of an application, you might use programmatic security to make access decisions. When administrative security is enabled and application server security is not disabled at the server level, Java EE applications security is enforced. When the security policy is specified for a Web resource, the Web container performs access control when the resource is requested by a Web client. The Web container challenges the Web client for authentication data if none is present according to the specified authentication method, ensures that the data constraints are met, and determines whether the authenticated user has the required security role. The Web security collaborator enforces role-based access control by using an access manager implementation. An access manager makes authorization decisions that are based on security policy derived from the deployment descriptor. An authenticated user principal can access the requested servlet or JSP file if the user principal has one of the required security roles. Servlets and JSP files can use the `HttpServletRequest` methods, `isUserInRole` and `getUserPrincipal`.

When administrative security and application security are enabled, and the application server level application security is not disabled, the EJB container enforces access control on EJB method invocation.

The authentication occurs regardless of whether method permission is defined for the specific EJB method. The EJB security collaborator enforces role-based access control by using an access manager implementation. An access manager makes authorization decisions that are based on security policy derived from the deployment descriptor. An authenticated user principal can access the requested EJB method if it has one of the required security roles. EJB code can use the `EJBContext` methods, `isCallerInRole` and `getCallerPrincipal`. Use the Java EE role-based access control to protect valuable business data from access by unauthorized users through the Internet and the intranet. Refer to *Securing Web applications using an assembly tool*, and *Securing enterprise bean applications*.

Role-based security

WebSphere Application Server extends the security, role-based access control to administrative resources including the JMX system management subsystem, user registries, and Java Naming and Directory Interface (JNDI) name space. WebSphere administrative subsystem defines four administrative security roles:

Monitor role

A monitor can view the configuration information and status but cannot make any changes.

Operator role

An operator can trigger run-time state changes, such as start an application server or stop an application but cannot make configuration changes.

Configurator role

A configurator can modify the configuration information but cannot change the state of the runtime.

Administrator role

An operator as well as a configurator, which additionally can modify sensitive security configuration and security policy such as setting server IDs and passwords, enable or disable administrative security and Java 2 security, and map users and groups to the administrator role.

iscadmins

The iscadmins role has administrator privileges for managing users and groups from within the administrative console only.

WebSphere Application Server defines two additional roles that are available when you use wsadmin scripting only.

Deployer

A deployer can perform both configuration actions and run-time operations on applications.

Adminsecuritymanager

An administrative security manager can map users to administrative roles. Also, when fine grained admin security is used, users granted this role can manage authorization groups.

Auditor

An auditor can view and modify the configuration settings for the security auditing subsystem.

A user with the configurator role can perform most administrative work including installing new applications and application servers. Certain configuration tasks exist that a configurator does not have sufficient authority to do when administrative security is enabled, including modifying a WebSphere Application Server identity and password, Lightweight Third-Party Authentication (LTPA) password and keys, and assigning users to administrative security roles. Those sensitive configuration tasks require the administrative role because the server ID is mapped to the administrator role.

Enable WebSphere Application Server administrative security to protect administrative subsystem integrity. Application server security can be selectively disabled if no sensitive information is available to protect. For securing administrative security, refer to “Authorizing access to administrative roles” on page 561 and Assigning users and groups to roles.

Java 2 security permissions

WebSphere Application Server uses the Java 2 security model to create a secure environment to run application code. Java 2 security provides a fine-grained and policy-based access control to protect system resources such as files, system properties, opening socket connections, loading libraries, and so on. The Java EE Version 1.4 specification defines a typical set of Java 2 security permissions that Web and EJB components expect to have. These permissions are shown in the following table.

Table 1. Java EE security permissions set for Web components

Security Permission	Target	Action
java.lang.RuntimePermission	loadLibrary	
java.lang.RuntimePermission	queuePrintJob	
java.net.SocketPermission	*	connect
java.io.FilePermission	*	read, write
java.util.PropertyPermission	*	read

Table 2. Java EE security permissions set for EJB components

Security Permission	Target	Action
java.lang.RuntimePermission	queuePrintJob	
java.net.SocketPermission	*	connect
java.util.PropertyPermission	*	read

The WebSphere Application Server Java 2 security default policies are based on the Java EE Version 1.4 specification. The specification grants Web components read and write file access permission to any file in the file system, which might be too broad. The WebSphere Application Server default policy gives Web components read and write permission to the subdirectory and the subtree where the Web module is installed. The default Java 2 security policies for all Java virtual machines and WebSphere Application Server processes are contained in the following policy files:

`\${java.home}/jre/lib/security/java.policy

This file is used as the default policy for the Java virtual machine (JVM).

`\${USER_INSTALL_ROOT}/properties/server.policy

This file is used as the default policy for all product server processes.

To simplify policy management, WebSphere Application Server policy is based on resource type rather than code base (location). The following files are the default policy files for a WebSphere Application Server subsystem. These policy files, which are an extension of the WebSphere Application Server runtime, are referred to as *Service Provider Programming Interfaces (SPI)*, and shared by multiple Java EE applications:

- *profile_root/config/cells/cell_name/nodes/node_name/spi.policy*
This file is used for embedded resources defined in the `resources.xml` file, such as the Java Message Service (JMS), JavaMail, and JDBC drivers.
- *profile_root/config/cells/cell_name/nodes/node_name/library.policy*
This file is used by the shared library that is defined by the WebSphere Application Server administrative console.
- *profile_root/config/cells/cell_name/nodes/node_name/app.policy*
This file is used as the default policy for Java EE applications.

In general, applications do not require more permissions to run than those recommended by the Java EE specification to be portable among various application servers. However, some applications might require more permissions. WebSphere Application Server supports the packaging of a `was.policy` file with each application to grant extra permissions to that application.

Note: Grant extra permissions to an application only after careful consideration because of the potential of compromising the system integrity.

Loading libraries into WebSphere Application Server does allow applications to leave the Java sandbox. WebSphere Application Server uses a permission filtering policy file to alert you when an application

installation fails because of additional permission requirements. For example, it is recommended that you not give the `java.lang.RuntimePermission` `exitVM` permission to an application so that application code cannot terminate WebSphere Application Server.

The filtering policy is defined by the filtermask in the `profile_root/config/cells/cell_name/filter.policy` file. Moreover, WebSphere Application Server also performs run-time permission filtering that is based on the run-time filtering policy to ensure that application code is not granted a permission that is considered harmful to system integrity.

Therefore, many applications developed for prior releases of WebSphere Application Server might not be Java 2 security ready. To quickly migrate those applications to the latest version of WebSphere Application Server, you might temporarily give those applications the `java.security.AllPermission` permission in the `was.policy` file. Test those applications to ensure that they run in an environment where Java 2 security is active. For example, identify which extra permissions, if any, are required, and grant only those permissions to a particular application. Not granting the `AllPermission` permission to applications can reduce the risk of compromising system integrity. For more information on migrating applications, refer to “Migrating Java 2 security policy” on page 29.

The WebSphere Application Server runtime uses Java 2 security to protect sensitive run-time functions. Applications that are granted the `AllPermission` permission not only have access to sensitive system resources, but also WebSphere Application Server run-time resources and can potentially cause damage to both. In cases where an application can be trusted as safe, WebSphere Application Server does support having Java 2 security disabled on a per application server basis. You can enforce Java 2 security by default in the administrative console and clear the Java 2 security flag to disable it at the particular application server.

When you specify the **Enable administrative security** and **Use Java 2 security to restrict application access to local resources** options on the Global security panel of the administrative console, the information and other sensitive configuration data, are stored in a set of XML configuration files. Both role-based access control and Java 2 security permission-based access control are employed to protect the integrity of the configuration data. The example uses configuration data protection to illustrate how system integrity is maintained.


Note: The **Enable global security** option in previous releases of WebSphere Application Server is the same as the **Enable administrative security** option in Version 7.0. Also, the **Enable Java 2 security** option in previous releases is the same as the **Use Java 2 security to restrict application access to local resources** option in Version 7.0.

- When Java 2 security is enforced, the application code cannot access the WebSphere Application Server run-time classes that manage the configuration data unless the code is granted the required WebSphere Application Server run-time permissions.
- When Java 2 security is enforced, application code cannot access the WebSphere Application Server configuration XML files unless the code is granted the required file read and write permission.
- The JMX administrative subsystem provides SOAP over HTTP or HTTPS and a RMI/IIOP remote interface to enable application programs to extract and to modify configuration files and data. When administrative security is enabled, an application program can modify the WebSphere Application Server configuration if the application program has presented valid authentication data and the security identity has the required security roles.
- If a user can disable Java 2 security, the user can also modify the WebSphere Application Server configuration, including the WebSphere Application Server security identity and authentication data with other sensitive data. Only users with the administrator security role can disable Java 2 security.
- Because WebSphere Application Server security identity is given to the administrator role, only users with the administrator role can disable administrative security, change server IDs and passwords, and map users and groups to administrative roles, and so on.

Other Runtime resources

Other WebSphere Application Server run-time resources are protected by a similar mechanism, as described previously. It is very important to enable WebSphere Application Server administrative security and to use Java 2 security to restrict application access to local resources. Java EE Specification defines several authentication methods for Web components: HTTP Basic Authentication, Form-Based Authentication, and HTTPS Client Certificate Authentication. When you use client certificate login, it is more convenient for the browser client if the Web resources have integral or confidential data constraint. If a browser uses HTTP to access the Web resource, the Web container automatically redirects the browser to the HTTPS port. The CSiv2 security protocol also supports client certificate authentication. You can also use SSL client authentication to set up secure communication among a selected set of servers based on a trust relationship.

If you start from the WebSphere Application Server plug-in at the Web server, you can configure SSL mutual authentication between it and the WebSphere Application Server HTTPS server. When using a certificate, you can restrict the WebSphere Application Server plug-in to communicate with only the selected two WebSphere Application Servers as shown in the following figure. Note that you can use self-signed certificates to reduce administration and cost.



For example, you want to restrict the HTTPS server in WebSphere Application Server **A** and in WebSphere Application Server **B** to accept secure socket connections only from the WebSphere Application Server plug-in **W**.

- To complete this task, you can generate three certificates using the IKEYMAN and the certificate management utilities. Also, you can use certificate **W** and trust certificate **A** and **B**. Configure the HTTPS server of WebSphere Application Server **A** to use certificate **A** and to trust certificate **W**.

Configure the HTTPS server of WebSphere Application Server **B** to use certificate **B** and to trust certificate **W**.

The trust relationship that is depicted in the previous figure is shown in the following table.

Server	Key	Trust
WebSphere Application Server plug-in	W	A, B
WebSphere Application Server A	A	W
WebSphere Application Server B	B	W

The WebSphere Application Server Deployment Manager is a central point of administration. System management commands are sent from the deployment manager to each individual application server. When administrative security is enabled, you can configure WebSphere Application Servers to require SSL and mutual authentication.

You might want to restrict WebSphere Application Server **A** so that it can communicate with WebSphere Application Server **C** only and WebSphere Application Server **B** can communicate with WebSphere Application Server **D** only. All WebSphere Application Servers must be able to communicate with WebSphere Application Server deployment manager **E**; therefore, when using self-signed certificates, you might configure the CSIV2 and SOAP/HTTPS Key and trust relationship, as shown in the following table.

Server	Key	Trust
WebSphere Application Server A	A	C, E
WebSphere Application Server B	B	D, E
WebSphere Application Server C	C	A, E
WebSphere Application Server D	D	B, E
WebSphere Application Server Deployment Manager E	E	A, B, C, D

When WebSphere Application Server is configured to use Lightweight Directory Access Protocol (LDAP) user registry, you also can configure SSL with mutual authentication between every application server and the LDAP server with self-signed certificates so that a password is not visible when it is passed from WebSphere Application Server to the LDAP server.

In this example, the node agent processes are not discussed. Each node agent must communicate with application servers on the same node and with the deployment manager. Node agents also must communicate with LDAP servers when configured to use an LDAP user registry. It is reasonable to let the deployment manager and the node agents use the same certificate. Suppose application server **A** and **C** are on the same computer node. The node agent on that node needs to have certificates **A** and **C** in its trust store.

WebSphere Application Server does not provide a registry configuration or management utility. In addition, it does not dictate the registry password policy. It is recommended that you use the password policy recommended by your registry, including the password length and expiration period.

Before securing your WebSphere Application Server environment, determine which versions of WebSphere Application Server you are using, review the WebSphere Application Server security architecture, and review each of the following topics:

- “Server and administrative security” on page 41
- “Authentication protocol support” on page 502
- “Common Secure Interoperability Version 2 features” on page 503
- “Identity assertion to the downstream server” on page 503
- “Selecting an authentication mechanism” on page 247
 - “Lightweight Third Party Authentication” on page 249
 - “Trust associations” on page 289

- “Single sign-on using LTPA cookies” on page 295
- “Selecting a registry or repository” on page 111
 - “Local operating system registries” on page 238
 - “Standalone Lightweight Directory Access Protocol registries” on page 242
- “Java 2 security” on page 44
 - “Java 2 security policy files” on page 48
- “Java Authentication and Authorization Service” on page 371
 - “Programmatic login” on page 393
- J2EE connector security
- “Access control exception” on page 52
 - “Role-based authorization” on page 516
 - “Administrative roles and naming service authorization” on page 510

Chapter 2. Task overview: Securing resources

WebSphere Application Server supports the Java Platform, Enterprise Edition (Java EE) model for creating, assembling, securing, and deploying applications. Applications are often created, assembled, and deployed in different phases and by different teams.

About this task

You can secure resources in a Java EE environment by following the required high-level steps. Consult the Java EE specifications for complete details.

- Set up and enable security. You must address several issues prior to authenticating users, authorizing access to resources, securing applications, and securing communications. These security issues include migration, interoperability, and installation. After installing WebSphere Application Server, you must determine the proper level of security that is needed for your environment. For more information, see Chapter 3, “Setting up and enabling security,” on page 15.
- Configure multiple domains. Security domains enable you to define multiple security configurations for use in your environment. For example, you can define different security (such as a different user registry) for user applications than for administrative applications. You can also define separate security configurations for user applications deployed to different servers and clusters. For more information, see Chapter 4, “Configuring multiple security domains,” on page 75
- Authenticate users. The process of authenticating users involves a user registry and an authentication mechanism. Optionally, you can define trust between WebSphere Application Server and a proxy server, configure single sign-on capability, and specify how to propagate security attributes between application servers. For more information, see Chapter 5, “Authenticating users,” on page 111.
- Authorize access to resources. WebSphere Application Server provides many different methods for authorizing accessing resources. For example, you can assign roles to users and configure a built-in or external authorization provider. For more information, see Chapter 6, “Authorizing access to resources,” on page 509.
- Secure communications. WebSphere Application Server provides several methods to secure communication between a server and a client. For more information, see Chapter 7, “Securing communications,” on page 589.
- Develop extensions to the WebSphere security infrastructure. WebSphere Application Server provides various plug points so that you can extend the security infrastructure. For more information, see Chapter 8, “Developing extensions to the WebSphere security infrastructure,” on page 731.
- Use the Auditing Facility to track and archive auditable events to ensure the integrity of your system. For more information, see Chapter 9, “Auditing the security infrastructure,” on page 839
- Secure various types of WebSphere applications. See **Securing WebSphere applications** for tasks involving developing, deploying, and administering secure applications, including Web applications, Web services, and many other types. This section highlights the security concerns and tasks that are specific to each type of application.
- Tune, harden, and maintain security configurations. After you have installed WebSphere Application Server, there are several considerations for tuning, strengthening, and maintaining your security configuration. For more information, see Chapter 13, “Tuning, hardening, and maintaining,” on page 1275.
- Troubleshoot security configurations. For more information, see Chapter 14, “Troubleshooting security configurations,” on page 1289.

Results

Your applications and production environment are secured.

Example

See [Security: Resources for learning](#) for more information on the WebSphere Application Server security architecture.

Chapter 3. Setting up and enabling security

You must address several issues prior to authenticating users, authorizing access to resources, securing applications, and securing communications. These security issues include migration, interoperability, and installation.

About this task

After installing WebSphere Application Server, you can determine the proper level of security that is needed for your environment. By default, administrative security is enabled and provides the authentication of users using the WebSphere administration functions, the use of Secure Sockets Layer (SSL), and the choice of user account repository.

The following information is covered in this section:

- Determine if any migration and interoperability issues might affect your installation. For more information, see “Migrating, coexisting, and interoperating – Security considerations.”
- Prepare your environment before and after installing WebSphere Application Server. For more information, see “Preparing for security at installation time” on page 34.
- Enable security for all your application servers or for specific application servers in your realm. For more information, see either “Enabling security” on page 37 or Chapter 4, “Configuring multiple security domains,” on page 75.

What to do next

After installing WebSphere Application Server and securing your environment, you must authenticate users. For more information, see Chapter 5, “Authenticating users,” on page 111.

Migrating, coexisting, and interoperating – Security considerations

Use this topic to migrate the security configuration of previous WebSphere Application Server releases and its applications to the new installation of WebSphere Application Server.

Before you begin

This information addresses the need to migrate your security configurations from a previous release of IBM WebSphere Application Server to WebSphere Application Server 7.0 or later. Complete the following steps to migrate your security configurations:

- If security is enabled in the previous release, obtain the administrative server ID and password of the previous release. This information is needed in order to run certain migration jobs.
- You can optionally disable security in the previous release before migrating the installation. No logon is required during the installation.

Use the First steps wizard to access and run the Migration wizard.

1. Start the First steps wizard by launching the `firststeps.bat` or the `firststeps.sh` file. The first steps file is located in the following directory:
 - `Linux` `AIX` `HP-UX` `Solaris` `UNIX` `./app_server_root/profiles/profile_name/firststeps/firststeps.sh`
 - `Windows` `app_server_root\profiles\profile_name\firststeps\firststeps.bat`
2. On the First steps wizard panel, click **Migration wizard**.
3. Follow the instructions provided in the First steps wizard to complete the migration.

Results

The security configuration of previous WebSphere Application Server releases and its applications are migrated to the new installation of WebSphere Application Server Version 7.0.

What to do next

If a custom user registry is used in the previous version, the migration process does not migrate the class files that are used by the standalone custom registry in the previous *app_server_root/classes* directory. Therefore, after migration, copy your custom user registry implementation classes to the *app_server_root/classes* directory.

If you upgrade from WebSphere Application Server, Version 5.x to WebSphere Application Server Version 7.0, the data that is associated with Version 5.x trust associations is not automatically migrated. To migrate trust associations, see “Migrating trust association interceptors” on page 22.

Interoperating with previous product versions

IBM WebSphere Application Server inter-operates with the previous product versions. Use this topic to configure this behavior.

Before you begin

The current release of the Application Server distinguishes the identities of the user who acts as an administrator, managing the Application Server environment, from the identity of the user that is used for authenticating between servers. In prior releases, the end user had to specify a server user ID and password as the user identity for authenticating between servers. In the current release of the Application Server, the server user ID is generated automatically and internally; however, the end user can specify that the server user ID and password not be automatically generated. This option is especially important in the case of a mixed-release cell, where the server user ID and password are specified in a down-level version of the Application Server. In such a scenario, the end user should opt out of automatically generating the server user ID and instead use the server user ID and password that is specified in the down-level version of the Application Server, in order to ensure backwards compatibility.

Interoperability is achieved only when the Lightweight Third Party Authentication (LTPA) authentication mechanism and a distributed user registry is used such as Lightweight Directory Access Protocol (LDAP) or a distributed Custom user registry. LocalOS on most platforms is not considered a distributed user registry (except on z/OS® within the z/OS environment). Also, the Simple WebSphere Authentication Mechanism (SWAM) cannot be used for interoperability as it does not contain credentials that can be forwarded outside of the existing process.

Note: SWAM was deprecated in WebSphere Application Server. Version 7.0 and will be removed in a future release.

1. Configure WebSphere Application Server Version 7.0 with the same distributed user registry (that is, LDAP or Custom) that is configured with the previous version. Make sure that the same LDAP user registry is shared by all of the product versions.
 - a. In the administrative console, select **Security > Global security**.
 - b. Choose an available Realm definition and click **Configure**.
 - c. Enter a **Primary administrative user name**. This identity is the user with administrative privileges that is defined in your local operating system. If you are not using the local OS as the user registry, select the **Server identity that is stored in the user repository**, enter the Server user ID, and the associated password. The user name is used to log on to the administrative console when administrative security is enabled. WebSphere Application Server Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 7.0, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

- d. When interoperating with Version 6.0.x or previous versions, you must select the Server identity that is stored in the user repository. Enter the **Server user id** and the associated **Password**.
2. Configure the LTPA authentication mechanism. Automatic generation of the LTPA keys should be disabled. If not, keys used by a previous release are lost. Export the current LTPA keys from WebSphere Application Server Version 7.0 and import them into the previous release.
 - a. In the administrative console select **Security > Global security**.
 - b. From Authentication mechanisms and expiration, click **LTPA**.
 - c. Click the **Key set groups** link , then click the key set group that displays in the Key set groups panel.
 - d. Clear the **Automatically generate keys** check box.
 - e. Click **OK**, then click **Authentication mechanisms and expiration** in the path at the top of the Key set groups panel.
 - f. Scroll down to the Cross-cell single sign-on section, and enter a password to use for encrypting the LTPA keys when adding them to the file.
 - g. Enter the password again to confirm the password.
 - h. Enter the **Fully qualified key file name** that contains the exported keys.
 - i. Click **Export keys**.
 - j. Follow the instructions provided in the previous release to import the exported LTPA keys into that configuration.
3. If you are using the default SSL configuration, extract all of the signer certificates from the WebSphere Application Server Version 7.0 common trust store. Otherwise, extract signers where necessary to import them into the previous release.
 - a. In the administrative console, click **Security > SSL certificate and key management**.
 - b. Click **Key stores and certificates**.
 - c. Click **CellDefaultTrustStore**.
 - d. Click **Signer certificates**.
 - e. Select one signer and click **Extract**.
 - f. Enter a unique path and filename for the signer. For example, /tmp/signer1.arm.
 - g. Click **OK**. Repeat for all of the signers in the trust store.
 - h. Check other trust stores for other signers that might need to be shared with the other server. Repeat steps e through h to extract the other signers.

You can also import a signer certificate, which is also called a certificate authority (CA) certificate, from a truststore on a non-z/OS platform server to a z/OS keyring. The z/OS keyring contains the signer certificates that originated on the non-z/OS platform server. For more information, see

4. Add the exported signers to DummyServerTrustFile.jks and DummyClientTrustFile.jks in the /etc directory of the back-level product version. If the previous release is not using the dummy certificate, the signer certificate(s) from the previous release must be extracted and added into the WebSphere Application Server Version 7.0 release to enable SSL connectivity in both directions.
 - a. Open the key management utility, iKeyman, for that product version.
 - b. Start ikeyman.bat or ikeyman.sh from the \${USER_INSTALL_ROOT}/bin directory.
 - c. Select **Key Database File > Open**.
 - d. Open \${USER_INSTALL_ROOT}/etc/DummyServerTrustFile.jks.
 - e. Enter WebAS for the password.

- f. Select **Add** and enter one of the files extracted in step 2. Continue until you have added all of the signers.
 - g. Repeat steps c through f for the `DummyClientTrustFile.jks` file.
5. Verify that the application uses the correct Java Naming and Directory Interface (JNDI) name and naming bootstrap port for performing a naming lookup.
 6. Stop and restart all of the servers.

Interoperating with a C++ common object request broker architecture client

WebSphere Application Server supports security in the CORBA C++ client to access-protected enterprise beans. If configured, C++ CORBA clients can access protected enterprise bean methods using a client certificate to achieve mutual authentication on WebSphere Application Server applications.

About this task

You can achieve interoperability of Security Authentication Service between the C++ Common Object Request Broker Architecture (CORBA) client and WebSphere Application Server using Common Secure Interoperability Version 2 (CSIv2) authentication protocol over Remote Method Invocation over the Internet Inter-ORB Protocol (RMI-IIOP). The CSIv2 security service protocol has authentication, attribute and transport layers. Among the three layers, transport authentication is conceptually simple, however, cryptographically based transport authentication is the strongest. WebSphere Application Server has implemented the transport authentication layer, so that C++ secure CORBA clients can use it effectively in making CORBA clients and protected enterprise bean resources work together.

Security authentication from non-Java based C++ client to enterprise beans. WebSphere Application Server supports security in the CORBA C++ client to access-protected enterprise beans. If configured, C++ CORBA clients can access protected enterprise bean methods using a client certificate to achieve mutual authentication on WebSphere Application Server applications.

To support the C++ CORBA client in accessing protected enterprise beans, complete the following steps:

- Create an environment file for the client, such as `current.env`. Set the variables presented in the following list in the file:

C++ security setting	Description
<code>client_protocol_password</code>	Specifies the password for the user ID.
<code>client_protocol_user</code>	Specifies the user ID to authenticate at the target server.
<code>security_sslKeyring</code>	Specifies the name of the RACF keyring for the client to use. The keyring must be defined under the user ID that is issuing the command to run the client.

- Point to the environment file using the fully qualified path name through the `WAS_CONFIG_FILE` environment variable. For example, in the `test.sh` test shell script, export:

```
/WebSphere/V6R0M0/DeploymentManager/profiles/default/config/cells
/PLEX1Network/nodes/PLEX1Manager/servers/dmgr
```

Some of the environment file terms are explained below:

default

profile name

PLEX1Network

cell name

PLEX1Manager

node name

dmgr server name

1. Obtain a valid certificate to represent the client and export its public key to the target enterprise bean server.

A valid certificate is needed to represent the C++ client. Request a certificate from the certificate authority (CA) or create a self-signed certificate for testing purposes.

Use the Key Management Utility from the Global Security Kit (GSKit) to extract the public key from the personal certificate and save it in the .arm format.

2. Prepare a truststore file for WebSphere Application Server.

Add the extracted client public key in the .arm file from the client to the server key truststore file. The server can now authenticate the client.

Note: This is done by invoking the Key Management Utility through `keyman.bat` or `keyman.sh` from WebSphere Application Server installation.

3. Configure WebSphere Application Server to support Secure Sockets Layer (SSL) as the authentication mechanism.

- a. Start the administrative console.

- b. Locate the application server that has the target enterprise bean deployed and configure it to use SSL client certificate authentication.

If it is a base installation, complete the following steps:

- 1) Click **Security > Global security**. Under RMI/IIOP security, click **CSlv2 inbound communications**. Select **Supported** for the Basic authentication and Client certificate authentication options. Leave the rest of the options as defaults.
- 2) Click **OK**.
- 3) Click **Security > Global security**. Under RMI/IIOP security, click **CSlv2 inbound communications** and verify that the **SSL-supported** option is selected under Transport.

If it is a Network Deployment setting, complete the following steps:

- 1) Click **Servers > Application Servers > server_name_where_the_EJB_resides**.
- 2) Under Security, click **Server security**.
- 3) Select the **RMI/IIOP security for this server overrides cell settings** option.
- 4) Under Additional properties, click **CSlv2 inbound communications**.
- 5) Select **Supported** for the Basic authentication and Client certificate authentication options. Leave the rest of the options as defaults.
- 6) Click **Servers > Application Servers > server_name_where_the_EJB_resides**.
- 7) Under Security, click **Server security**.
- 8) Under Additional properties, click **CSlv2 inbound communications**.
- 9) Verify that the **SSL-Supported** option is selected.

- c. Restart the application server.

The WebSphere Application Server is ready to take a C++ CORBA security client and a mutually authenticated server and client by using SSL in the transport layer.

4. Configure the C++ CORBA client to use a certificate in performing the mutual authentication.

Client users are accustomed to using property files in their applications because they are helpful in specifying configuration settings. The following list presents important C++ security settings:

C++ security setting	Description
<code>com.ibm.CORBA.bootstrapHostName=ricebella.austin.ibm.com</code>	Specifies the target host name.
<code>com.ibm.CORBA.securityEnabled=yes</code>	Enables security.
<code>com.ibm.CSI.performTLClientAuthenticationSupported=yes</code>	Ensures client is supporting mutual authentication by certificate
<code>com.ibm.ssl.keyFile=C:/ricebella/etc/DummyKeyRingFile.KDB</code>	Specifies which key database file to use.

C++ security setting	Description
com.ibm.ssl.keyPassword=WebAS	Specifies the password for opening the key database file. WebSphere Application Server supports a utility called PasswordEncode4cpp to encode the plain password.
com.ibm.CORBA.translationEnabled=1	Enables the valueType conversion.

To use the property files in running a C++ client, an environment variable WASPROPS, is used to indicate where a property file or a list of property files exists.

For the complete set of C++ client properties, see the sample property file `sccClient.props`, which is shipped with the product located in the `app_server_root/profiles/profile_name/etc` directory.

Migrating custom user registries

If you built your own custom user registry, consider the migration items listed below. If you have a custom user registry that was provided by a Security Solution Provider, you must contact that provider to ensure that you have the correct version of their custom user registry to support WebSphere Application Server.

Before you begin

In WebSphere Application Server, in addition to the UserRegistry interface, the custom user registry requires the Result object to handle user and group information. This file is already provided in the package and you are expected to use it for the `getUsers`, `getGroups`, and the `getUsersForGroup` methods.

You cannot use other WebSphere Application Server components, for example, data sources, to initialize the custom registry because other components, like the containers, are initialized after security and are not available during the registry initialization. A custom registry implementation is a pure custom implementation, independent of other WebSphere Application Server components.

The `getCallerPrincipal` enterprise bean method and the `getUserPrincipal` and `getRemoteUser` servlet methods return the security name instead of the display name. For more information, see the API documentation.

If the migration tool is used to migrate the WebSphere Application Server Version 5 configuration to WebSphere Application Server Version 6.0.x and later, this migration does not change your existing code. Because the WebSphere Application Server Version 5 custom registry works in WebSphere Application Server Version 6.0.x and later without any changes to the implementation, except when using data sources, you can use the Version 5-based custom registry after the migration without modifying the code.

Copy the files to the cell and to each of the node class paths.

In WebSphere Application Server Version 6.0.x and later, a case-insensitive authorization can occur when using an enabled custom user registry.

Setting this flag does not have any effect on the user names or passwords. Only the unique IDs that are returned from the registry are changed to lower-case before comparing them with the information in the authorization table, which is also converted to lowercase during runtime.

Before proceeding, look at the UserRegistry interface. See “Developing standalone custom registries” on page 731 for a description of each of these methods in detail.

About this task

The following steps go through all the changes that are required to move your WebSphere Application Server Version 4.x custom user registry that implemented the old `com.ibm.websphere.security.CustomRegistry` interface to the `com.ibm.websphere.security.UserRegistry` interface.

Note: The sample implementation file is used as an **example** when describing the following steps.

1. Change your implementation to `UserRegistry` instead of `CustomRegistry`. Change:

```
public class FileRegistrySample implements CustomRegistry
```

to:

```
public class FileRegistrySample implements UserRegistry
```

2. Create the `java.rmi.RemoteException` exception in the constructors:

```
public FileRegistrySample() throws java.rmi.RemoteException
```

3. Change the `mapCertificate` method to take a certificate chain instead of a single certificate. Change

```
public String mapCertificate(X509Certificate cert)
```

to:

```
public String mapCertificate(X509Certificate[] cert)
```

Having a certificate chain gives you the flexibility to act on the chain instead of one certificate. If you are interested only in the first certificate, take the first certificate in the chain before processing. In WebSphere Application Server Version 6.0.x and later, the `mapCertificate` method is called to map the user in a certificate to a valid user in the registry when certificates are used for authentication by the Web or the Java clients.

4. Remove the `getUsers` method.
5. Change the signature of the `getUsers(String)` method to return a `Result` object and accept an additional parameter (`int`). Change:

```
public List getUsers(String pattern)
```

to:

```
public Result getUsers(String pattern, int limit)
```

In your implementation, construct the `Result` object from the list of the users that is obtained from the user registry (whose number is limited to the value of the `limit` parameter) and call the `setHasMore` method on the `Result` object if the total number of users in the registry exceeds the `limit` value.

6. Change the signature of the `getUsersForGroup(String)` method to return a `Result` object and accept an additional parameter (`int`) and throw a new exception called `NotImplementedException` exception. Change the following code:

```
public List getUsersForGroup(String groupName)
    throws CustomRegistryException,
           EntryNotFoundException {
```

to:

```
public Result getUsersForGroup(String groupSecurityName, int limit)
    throws NotImplementedException,
           EntryNotFoundException,
           CustomRegistryException {
```

In WebSphere Application Server Version 6.0.x and later, this method is not called directly by the WebSphere Application Server security component. However, other components of WebSphere Application Server, like the WebSphere Business Integration Server Foundation process choreographer, use this method when staff assignments are modeled using groups. Because this implementation is supported in WebSphere Application Server Version 6.0.x and later, it is recommended that you change the implementation similar to the `getUsers` method as explained in step 5.

7. Remove the `getUniqueUserIds(String)` method.
8. Remove the `getGroups` method.
9. Change the signature of the `getGroups(String)` method to return a `Result` object and accept an additional parameter (`int`). Change the following code:

```
public List getGroups(String pattern)
```

to:

```
public Result getGroups(String pattern, int limit)
```

In your implementation, construct the Result object from the list of the groups that is obtained from the user registry whose number is limited to the value of the limit parameter. Call the `setHasMore` method on the Result object if the total number of groups in the registry exceeds the limit value.

10. Add the `createCredential` method. This method is not called at this time, so return as `null`.

```
public com.ibm.websphere.security.cred.WSCredential  
    createCredential(String userSecurityName)  
        throws CustomRegistryException,  
            NotImplementedException,  
            EntryNotFoundException {  
    return null;  
}
```

The first and second lines of the previous code example are split onto two lines for illustrative purposes only.

11. To build the WebSphere Application Server Version 6.0.x and later implementation, make sure you have the `com.ibm.ws.runtime.jar` file in your class path.

To set the files in your class path, use the following code as a sample and substitute your environment values for the variables that are used in the example:

```
app_server_root/java/bin/javac -classpath  
app_server_root/plugins/com.ibm.ws.runtime.jar FileRegistrySample.java
```

Type the previous lines as one continuous line.

To build the WebSphere Application Server Version 5 custom registry (CustomRegistry) in WebSphere Application Server Version 6.0.x and later, only the `com.ibm.ws.runtime.jar` file is required.

12. Copy the implementation classes to the product class path.

The `%install_root%/lib/ext` directory is the preferred location.

Make sure that you copy these classes to the cell and all the nodes. Without the classes in each of the node class paths, the nodes and the application servers in those nodes cannot start when security is on.

13. Use the administrative console to set up the custom registry.

Follow the instructions in “Configuring standalone custom registries” on page 142 to set up the custom registry, including the **ignore case for authorization** option. Make sure that you add the `WAS_UseDisplayName` properties if required.

Results

WebSphere Application Server Version 4.x based custom user registry that implemented the old `com.ibm.websphere.security.CustomRegistry` interface is migrated to the `com.ibm.websphere.security.UserRegistry` interface.

What to do next

If you are enabling security, see “Enabling security” on page 37 to complete the remaining steps. When completed, save the configuration and restart all the servers. Try accessing some Java Platform, Enterprise Edition (Java EE) resources to verify that the custom registry migration is successful.

Migrating trust association interceptors

Use this topic to manually migrate trust associations.

Before you begin

Note: Data sources are not supported for use within a Trust Association Interceptor (TAI). Data sources are intended for use within J2EE applications and designed to operate within the EJB and Web containers. Trust Association Interceptors do not run within a container, and while data sources may function in the TAI environment, they are untested and not guaranteed to function properly.

The following topics are addressed in this document:

- Changes to the product-provided trust association interceptors
- Migrating product-provided trust association interceptors
- Changes to the custom trust association interceptors
- Migrating custom trust association interceptors

Changes to the product-provided trust association interceptors

For the product-provided implementation for the WebSEAL server, a new optional `com.ibm.websphere.security.webseal.ignoreProxy` property is added. If this property is set to `true` or `yes`, the implementation does not check for the proxy host names and the proxy ports to match any of the host names and ports that are listed in the `com.ibm.websphere.security.webseal.hostnames` and the `com.ibm.websphere.security.webseal.ports` property respectively. For example, if the VIA header contains the following information:

```
HTTP/1.1 Fred (Proxy), 1.1 Sam (Apache/1.1),  
HTTP/1.1 webseal1:7002, 1.1 webseal2:7001
```

and the `com.ibm.websphere.security.webseal.ignoreProxy` property is set to `true` or `yes`, the host name `Fred`, is not used when matching the host names. By default, this property is not set, which implies that any proxy host names and ports that are expected in the VIA header are listed in the host names and the ports properties to satisfy the `isTargetInterceptor` method.

The previous VIA header information was split onto two lines for illustrative purposes only.

For more information about the `com.ibm.websphere.security.webseal.ignoreProxy` property, see the article in the information center on configuring single signon using trust association interceptor ++.

Migrating product-provided trust association interceptors

The properties that are located in the `webseal.properties` and `trustedserver.properties` files are not migrated from previous versions of WebSphere Application Server. You must migrate the appropriate properties to WebSphere Application Server Version 6.0.x using the trust association panels in the administrative console. For more information, see [Configuring trust association interceptors](#).

Changes to the custom trust association interceptors

If the custom interceptor extends the `com.ibm.websphere.security.WebSphereBaseTrustAssociationInterceptor` property, implement the following new method to initialize the interceptor:

```
public int init (java.util.Properties props);
```

WebSphere Application Server checks the return status before using the trust association implementation. Zero (0) is the default value for indicating that the interceptor is successfully initialized.

However, if a previous implementation of the trust association interceptor returns a different error status, you can either change your implementation to match the expectations or make one of the following changes:

Method 1:

Add the `com.ibm.websphere.security.trustassociation.initStatus` property in the trust association interceptor custom properties. Set the property to the value that indicates the interceptor is successfully initialized. All of the other possible values imply failure. In case of failure, the corresponding trust association interceptor is not used.

Method 2:

Add the `com.ibm.websphere.security.trustassociation.ignoreInitStatus` property in the trust association interceptor custom properties. Set the value of this property to `true`, which tells WebSphere Application Server to ignore the status of this method. If you add this property to the custom properties, WebSphere Application Server does not check the return status, which is similar to previous versions of WebSphere Application Server.

The public `int init (java.util.Properties props)` method replaces the public `int init (String propsFile)` method.

The `init(Properties)` method accepts a `java.util.Properties` object, which contains the set of properties that is required to initialize the interceptor. All of the properties set for an interceptor are sent to this method. The interceptor can then use these properties to initialize itself. For example, in the product-provided implementation for the WebSEAL server, this method reads the hosts and ports so that a request coming in can be verified to come from trusted hosts and ports. A return value of Zero (0) implies that the interceptor initialization is successful. Any other value implies that the initialization is not successful and the interceptor is not used.

The `init(String)` method still works if you want to use it instead of implementing the `init(Properties)` method. The only requirement is that you enter the file name containing the custom trust association properties using the **Custom Properties** link of the interceptor in the administrative console or by using scripts. You can enter the property using either of the following methods. The first method is used for backward compatibility with previous versions of WebSphere Application Server.

Method 1:

The same property names used in the previous release are used to obtain the file name. The file name is obtained by concatenating `.config` to the `com.ibm.websphere.security.trustassociation.types` property value. If the `myTAI.properties` file is located in the `app_server_root/properties` directory, set the following properties:

- `com.ibm.websphere.security.trustassociation.types = myTAItype`
- `com.ibm.websphere.security.trustassociation.myTAItype.config = app_server_root/properties/myTAI.properties`

Method 2:

You can set the `com.ibm.websphere.security.trustassociation.initPropsFile` property in the trust association custom properties to the location of the file. For example, set the following property:

```
com.ibm.websphere.security.trustassociation.initPropsFile=  
app_server_root/properties/myTAI.properties
```

The previous line of code is split into two lines for illustrative purposes only. Type as one continuous line.

In a Network Deployment installation, where the location of the file name can vary for different nodes, use the variable `install_root` to refer to the WebSphere Application Server installation directory.

However, it is highly recommended that your implementation be changed to implement the `init(Properties)` method instead of relying on the `init (String propsfile)` method.

Migrating custom trust association interceptors

The trust associations from previous versions of WebSphere Application Server are not automatically migrated to WebSphere Application Server Version 6.0.x and later. You can manually migrate these trust associations using the following steps:

1. Recompile the implementation file, if necessary.

For more information, refer to the "Changes to the custom trust association interceptors" section previously discussed in this document.

To recompile the implementation file, type the following code:

```
%WAS_HOME%/java/bin/javac -classpath %WAS_HOME%/plugins/com.ibm.ws.runtime.jar;  
%WAS_HOME%/lib/j2ee.jar your_implementation_file.java
```

The previous line of code is broken into two lines for illustrative purposes only. Type the code as one continuous line.

2. Copy the custom trust association interceptor class files to a location in your product class path. Copy these class files into the %WAS_HOME%/lib/ext directory.

You must copy this class file into the class path of each node and cell.

3. Start WebSphere Application Server.
4. Enable security to use the trust association interceptor. The properties that are located in your custom trust association properties file and in the `trustedserver.properties` file are not migrated from previous versions of WebSphere Application Server. You must migrate the appropriate properties to WebSphere Application Server Version 6.0.x and later using the trust association panels in the administrative console.

For more information, see *Configuring trust association interceptors*.

Migrating Common Object Request Broker Architecture programmatic login to Java Authentication and Authorization Service (CORBA and JAAS)

Use this topic as an example of how to perform programmatic login using the CORBA-based programmatic login APIs.

Before you begin

This document outlines the deprecated Common Object Request Broker Architecture (CORBA) programmatic login APIs and the alternatives that are provided by JAAS. WebSphere Application Server fully supports the Java Authentication and Authorization Service (JAAS) as programmatic login application programming interfaces (API). Refer to the *Securing applications and their environment* PDF for more details on JAAS support.

The following list includes the deprecated CORBA programmatic login APIs.

- `${user.install.root}/installedApps/sampleApp.ear/default_app.war/WEB-INF/classes/LoginHelper.java`.
- `${user.install.root}/installedApps/sampleApp.ear/default_app.war/WEB-INF/classes/ServerSideAuthenticator.java`.
- **org.omg.SecurityLevel2.Credentials**. This API is included with the product, but it is not recommended that you use the API.

The APIs that are provided in WebSphere Application Server are a combination of standard JAAS APIs and a product implementation of standard JAAS interfaces.

The following information is only a summary; refer to the JAAS documentation for your platform located at: <http://www.ibm.com/developerworks/java/jdk/security/>.

- Programmatic login APIs:
 - `javax.security.auth.login.LoginContext`
 - `javax.security.auth.callback.CallbackHandler` interface: The WebSphere Application Server product provides the following implementation of the `javax.security.auth.callback.CallbackHandler` interface:
com.ibm.websphere.security.auth.callback.WSCallbackHandlerImpl
Provides a non-prompt `CallbackHandler` handler when the application pushes basic

authentication data (user ID, password, and security realm) or token data to product login modules. This API is recommended for server-side login.

com.ibm.websphere.security.auth.callback.WSGUICallbackHandlerImpl

Provides a login prompt CallbackHandler handler to gather basic authentication data (user ID, password, and security realm). This API is recommended for client-side login.

If this API is used on the server side, the server is blocked for input.

com.ibm.websphere.security.auth.callback.WSStdinCallbackHandlerImpl

Provides a stdin login prompt CallbackHandler handler to gather basic authentication data (user ID, password, and security realm). This API is recommended for client-side login.

Note: If this API is used on the server side, the server is blocked for input.

- javax.security.auth.callback.Callback interface:

javax.security.auth.callback.NameCallback

Provided by JAAS to pass the user name to the LoginModules interface.

javax.security.auth.callback.PasswordCallback

Provided by JAAS to pass the password to the LoginModules interface.

com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl

Provided by the product to perform a token-based login. With this API, an application can pass a token-byte array to the LoginModules interface.

- javax.security.auth.spi.LoginModule interface

WebSphere Application Server provides a LoginModules implementation for client and server-side login. Refer to the *Securing applications and their environment* PDF for details.

- javax.security.Subject:

com.ibm.websphere.security.auth.WSSubject

An extension provided by the product to invoke remote J2EE resources using the credentials in the javax.security.Subject

com.ibm.websphere.security.cred.WSCredential

After a successful JAAS login with the WebSphere Application Server LoginModules interfaces, a com.ibm.websphere.security.cred.WSCredential credential is created and stored in the Subject.

com.ibm.websphere.security.auth.WSPrincipal

An authenticated principal that is created and stored in a Subject that is authenticated by the WebSphere Application Server LoginModules interface.

1. Use the following as an example of how to perform programmatic login using the CORBA-based programmatic login APIs: The CORBA-based programmatic login APIs are replaced by JAAS login.

Note: The LoginHelper application programming interface (API) that is used in the following example is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release. It is recommended that you use the JAAS programmatic login APIs that are shown in the next step.

```
public class TestClient {
    ...
    private void performLogin() {
        // Get the ID and password of the user.
        String userid = customGetUserid();
        String password = customGetPassword();

        // Create a new security context to hold authentication data.
        LoginHelper loginHelper = new LoginHelper();
        try {
            // Provide the ID and password of the user for authentication.
            org.omg.SecurityLevel2.Credentials credentials =
                loginHelper.login(userid, password);

            // Use the new credentials for all future invocations.
            loginHelper.setInvocationCredentials(credentials);
            // Retrieve the name of the user from the credentials
            // so we can tell the user that login succeeded.
        }
    }
}
```



```

String username = loginHelper.getUserName(credentials);
System.out.println("Security context set for user: "+username);
} catch (org.omg.SecurityLevel2.LoginFailed e) {
// Handle the LoginFailed exception.
}
}
...
}

```

2. Use the following example to migrate the CORBA-based programmatic login APIs to the JAAS programmatic login APIs.

The following example assumes that the application code is granted for the required Java 2 security permissions. For more information, see the *Securing applications and their environment* PDF and the JAAS documentation located at <http://www.ibm.com/developerworks/java/jdk/security/>.

```

public class TestClient {
...
private void performLogin() {
// Create a new JAAS LoginContext.
javax.security.auth.login.LoginContext lc = null;

try {
// Use GUI prompt to gather the BasicAuth data.
lc = new javax.security.auth.login.LoginContext("WSLogin",
new com.ibm.websphere.security.auth.callback.WSGUICallbackHandlerImpl());

// create a LoginContext and specify a CallbackHandler implementation
// CallbackHandler implementation determine how authentication data is collected
// in this case, the authentication date is collected by login prompt
// and pass to the authentication mechanism implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: failed to instantiate a LoginContext and the exception: "
+ e.getMessage());
e.printStackTrace();

// may be javax.security.auth.AuthPermission "createLoginContext" is not granted
// to the application, or the JAAS Login Configuration is not defined.
}

if (lc != null)
try {
lc.login(); // perform login
javax.security.auth.Subject s = lc.getSubject();
// get the authenticated subject

// Invoke a J2EE resources using the authenticated subject
com.ibm.websphere.security.auth.WSSubject.doAs(s,
new java.security.PrivilegedAction() {
public Object run() {
try {
bankAccount.deposit(100.00); // where bankAccount is an protected EJB
} catch (Exception e) {
System.out.println("ERROR: error while accessing EJB resource, exception: "
+ e.getMessage());
e.printStackTrace();
}
return null;
}
}
);

// Retrieve the name of the principal from the Subject
// so we can tell the user that login succeeded,
// should only be one WSPPrincipal.
java.util.Set ps =
s.getPrincipals(com.ibm.websphere.security.auth.WSPPrincipal.class);
java.util.Iterator it = ps.iterator();
while (it.hasNext()) {
com.ibm.websphere.security.auth.WSPPrincipal p =
(com.ibm.websphere.security.auth.WSPPrincipal) it.next();
System.out.println("Principal: " + p.getName());
}
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: login failed with exception: " + e.getMessage());
}
}
}

```

```
e.printStackTrace();

// login failed, might want to provide relogin logic
}
}
...
}
```

Migrating from the CustomLoginServlet class to servlet filters

Use this topic to allow migration in an application that uses form-based login and servlet filters without the use of the CustomLoginServlet class.

Before you begin

The CustomLoginServlet class is deprecated in WebSphere Application Server Version 5. Those applications using the CustomLoginServlet class to perform authentication now need to use form-based login. Using the form-based login mechanism, you can control the look and feel of the login screen. In form-based login, a login page is specified and displays when retrieving the user ID and password information. You also can specify an error page that displays when authentication fails.

If login and error pages are not enough to implement the CustomLoginServlet class, use servlet filters. Servlet filters can dynamically intercept requests and responses to transform or use the information that is contained in the requests or responses. One or more servlet filters attach to a servlet or a group of servlets. Servlet filters also can attach to JavaServer Pages (JSP) files and HTML pages. All the attached servlet filters are called before invoking the servlet.

Both form-based login and servlet filters are supported by any Servlet 2.3 specification-compliant Web container. A form login servlet performs the authentication and servlet filters can perform additional authentication, auditing, or logging tasks.

To perform pre-login and post-login actions using servlet filters, configure these servlet filters for either form login page or for /j_security_check URL. The j_security_check is posted by the form login page with the j_username parameter that contains the user name and the j_password parameter that contains the password. A servlet filter can use user name and password information to perform more authentication or meet other special needs.

1. Develop a form login page and error page for the application.
Refer to the *Securing applications and their environment* PDF for details.
2. Configure the form login page and the error page for the application as described in .
Refer to the *Securing applications and their environment* PDF for details.
3. Develop servlet filters if additional processing is required before and after form login authentication.
Refer to the *Securing applications and their environment* PDF for details.
4. Configure the servlet filters that are developed in the previous step for either the form login page URL or for the /j_security_check URL. Use an assembly tool or development tools like Rational® Application Developer to configure filters. After configuring the servlet filters, the web-xml file contains two stanzas. The first stanza contains the servlet filter configuration, the servlet filter, and its implementation class. The second stanza contains the filter mapping section and a mapping of the servlet filter to the URL. For more information, see the *Securing applications and their environment* PDF.

Results

This migration results in an application that uses form-based login and servlet filters without the use of the CustomLoginServlet class.

What to do next

The new application uses form-based login and servlet filters to replace the CustomLoginServlet class. Servlet filters also are used to perform additional authentication, auditing, and logging.

Migrating Java 2 security policy

Use this topic for guidance pertaining to migrating Java 2 security policy.

About this task

Previous WebSphere Application Server releases

WebSphere Application Server uses the Java 2 security manager in the server runtime to prevent enterprise applications from calling the System.exit and the System.setSecurityManager methods. These two Java application programming interfaces (API) have undesirable consequences if called by enterprise applications. The System.exit API, for example, causes the Java virtual machine (application server process) to exit prematurely, which is not a beneficial operation for an application server.

To support Java 2 security properly, all the server runtime must be marked as privileged (with doPrivileged API calls inserted in the correct places), and identify the default permission sets or policy. Application code is not privileged and subject to the permissions that are defined in the policy files. The doPrivileged instrumentation is important and necessary to support Java 2 security. Without it, the application code must be granted the permissions that are required by the server runtime. This situation is due to the design and algorithm that is used by Java 2 security to enforce permission checks. Refer to the Java 2 security check permission algorithm.

The following two permissions are enforced by the Java 2 security manager (hard coded) for WebSphere Application Server:

- java.lang.RuntimePermission(exitVM)
- java.lang.RuntimePermission(setSecurityManager)

Application code is denied access to these permissions regardless of what is in the Java 2 security policy. However, the server runtime is granted these permissions. All the other permission checks are not enforced.

Only two permissions are supported:

- java.net.SocketPermission
- java.net.NetPermission

However, not all the product server runtime is properly marked as privileged. You must grant the application code all the other permissions besides the two listed previously or the enterprise application can potentially fail to run. This Java 2 security policy for enterprise applications is liberal.

What changed

Java 2 Security is fully supported in WebSphere Application Server, which means that all permissions are enforced. The default Java 2 security policy for an enterprise application is the recommended permission set defined by the Java Platform, Enterprise Edition (Java EE) Version 1.4 specification. Refer to the *profile_root/config/cells/cell_name/nodes/node_name/app.policy* file for the default Java 2 security policy that is granted to enterprise applications. This policy is a much more stringent compared to previous releases.

All policy is declarative. The product security manager honors all policy that is declared in the policy files. There is an exception to this rule: enterprise applications are denied access to permissions that are declared in the *profile_root/config/cells/cell_name/filter.policy* file.

Note: The default Java 2 security policy for enterprise applications is much more stringent and all the permissions are enforced in WebSphere Application Server Version 6.0.x and later. The security policy might fail because the application code does not have the necessary permissions granted where system resources, such as file I/O, can be programmatically accessed and are now subject to the permission checking.

In application code, do not use the `setSecurityManager` permission to set a security manager. When an application uses the `setSecurityManager` permission, there is a conflict with the internal security manager within WebSphere Application Server. If you must set a security manager in an application for RMI purposes, you also must enable the **Use Java 2 security to restrict application access to local resources** option on the Global security page within the WebSphere Application Server administrative console. WebSphere Application Server then registers a security manager. The application code can verify that this security manager is registered by using `System.getSecurityManager()` application programming interface (API).

Migrating system properties

The following system properties are used in previous releases in relation to Java 2 security:

- **java.security.policy.** The absolute path of the policy file (action required). This system property contains both system permissions (permissions granted to the Java virtual machine (JVM) and the product server runtime) and enterprise application permissions. Migrate the Java 2 security policy of the enterprise application to WebSphere Application Server Version 6.0.x. For Java 2 security policy migration, see the steps for migrating Java 2 security policy.
- **enableJava2Security.** Used to enable Java 2 security enforcement (no action required). This system property is deprecated; a flag in the WebSphere configuration application programming interface (API) is used to control whether to enable Java 2 security. Enable this option through the administrative console.
- **was.home.** Expanded to the installation directory of WebSphere Application Server (action might be required). This system property is deprecated; superseded by the `${user.install.root}` and `${was.install.root}` properties. If the directory contains instance-specific data then `${user.install.root}` is used; otherwise `${was.install.root}` is used. Use these properties interchangeably for the WebSphere Application Server or the Network Deployment environments. See the steps for migrating Java 2 security policy.

Migrating the Java 2 Security Policy

No easy way exists to migrate the Java policy file to WebSphere Application Server Version 6.0.x and later automatically because of a mixture of system permissions and application permissions in the same policy file. Manually copy the Java 2 security policy for enterprise applications to a `was.policy` or `app.policy` file. However, migrating the Java 2 security policy to a `was.policy` file is preferable because symbols or relative code base is used instead of an absolute code base. This process has many advantages. Grant the permissions that are defined in the `was.policy` to the specific enterprise application only, while permissions in the `app.policy` file apply to all the enterprise applications that run on the node where the `app.policy` file belongs.

Refer to the *Securing applications and their environment* PDF for more details on policy management.

The following example illustrates the migration of a Java 2 security policy from a previous release. The contents include the Java 2 security policy file for the `app1.ear` enterprise application and the system permissions, which are permissions that are granted to the Java virtual machine (JVM) and the product server runtime.

The default location for the Java 2 security policy file is `profile_root/properties/java.policy`. Default permissions are omitted for clarity:

```
// For product Samples
grant codeBase "file:${app_server_root}/installedApps/app1.ear/-" {
    permission java.security.SecurityPermission "printIdentity";
    permission java.io.FilePermission "${app_server_root}${/}temp${/}somefile.txt",
        "read";
};
```

For clarity of illustration, all the permissions are migrated as the application level permissions in this example. However, you can grant permissions at a more granular level at the component level (Web, enterprise beans, connector or utility Java archive (JAR) component level) or you can grant permissions to a particular component.

1. Ensure that Java 2 security is disabled on the application server.
2. Create a new `was.policy` file, if the file is not present, or update the `was.policy` file for migrated applications in the configuration repository with the following contents:

```
grant codeBase "file:${application}" {
    permission java.security.SecurityPermission "printIdentity";
    permission java.io.FilePermission "
        ${user.install.root}${/}temp${/}somefile.txt", "read";
};
```

The third and fourth lines in the previous code sample are presented on two lines for illustrative purposes only.

The `was.policy` file is located in the `profile_root/config/cells/cell_name/applications/app.ear/deployments/app/META-INF/` directory.

3. Use an assembly tool to attach the `was.policy` file to the enterprise archive (EAR) file. You also can use an assembly tool to validate the contents of the `was.policy` file. For more information, see the *Securing applications and their environment* PDF.
4. Validate that the enterprise application does not require additional permissions to the migrated Java 2 security permissions and the default permissions set declared in the `${user.install.root}/config/cells/cell_name/nodes/node_name/app.policy` file. This validation requires code review, code inspection, application documentation review, and sandbox testing of migrated enterprise applications with Java 2 security enabled in a preproduction environment. Refer to developer kit APIs protected by Java 2 security for information about which APIs are protected by Java 2 security. If you use third-party libraries, consult the vendor documentation for APIs that are protected by Java 2 security. Verify that the application is granted all the required permissions, or it might fail to run when Java 2 security is enabled.
5. Perform preproduction testing of the migrated enterprise application with Java 2 security enabled. Enable trace for the WebSphere Application Server Java 2 security manager in a preproduction testing environment with the following trace string: `com.ibm.ws.security.core.SecurityManager=all=enabled`. This trace function can be helpful in debugging the `AccessControlException` exception that is created when an application is not granted the required permission or some system code is not properly marked as privileged. The trace dumps the stack trace and permissions that are granted to the classes on the call stack when the exception is created.

For more information, see the *Securing applications and their environment* PDF.

Note: Because the Java 2 security policy is much more stringent compared with previous releases, the administrator or deployer must review their enterprise applications to see if extra permissions are required before enabling Java 2 security. If the enterprise applications are not granted the required permissions, they fail to run.

Migrating with Tivoli Access Manager for authentication enabled on a single node

When Tivoli® Access Manager security is configured for your existing environment and security is enabled for a single node, you can migrate to WebSphere Application Server, Version 7.0.

Before you begin

Your profiles must be migrated using the migration tools to migrate product configurations.

Note: Do not restart the WebSphere Application Server Version 7.0 server until after performing the following procedure. The migration tools omit some files that enable the server to start correctly.

About this task

After migrating your profiles additional steps are required when Tivoli Access Manager security is configured.

1. Copy the following files from the existing directory to the same directory for Version 7.0.

```
%WAS_HOME%\java\jre\PDPerm.properties
%WAS_HOME%\java\jre\lib\security\PdPerm.ks
%WAS_HOME%\java\jre\PolicyDirector\PDCA.ks
```

2. Edit the PD.properties file, and change the following configuration settings:

```
appsvr-plcysvrs=null\:\:1
config_type=standalone
```

Make the appropriate changes to point to your Tivoli Access Manager Policy Server, for example:

```
appsvr-plcysvrs=pdmgrd.test.gc.au.ibm.com\:\:7135\:\:1
config_type=full
```

3. Edit the PdPerm.properties file, and change all path names to the correct path name. Change the following configuration settings:

```
pdvar-home=C:\:\Program
Files\WebSphere\AppServer\java\jre\PolicyDirector
baseGroup.PDJv1dugong-v2dugongMessageFileHandler.fileName=C:\:\Program
Files\WebSphere\AppServer\java\jre\PolicyDirector\log/msg__v1dugong-v2dugong.log

pdcert-url=file\:\:c\:\:progra~1\WebSphere/AppServer/java/jre/lib/security/PdPerm.ks

baseGroup.PDJv1dugong-v2dugongTraceFileHandler.fileName=C:\:\Program
Files\WebSphere\AppServer\java\jre\PolicyDirector\log/trace__v1dugong-v2dugong.log

pd-home=C:\:\Program Files\WebSphere\AppServer\java\jre\PolicyDirector

java-home=C:\:\Program Files\WebSphere\AppServer\java\jre
```

What to do next

Also see the migration information with Tivoli Access Manager for authentication that is enabled on multiple nodes with security enabled.

Migrating with Tivoli Access Manager for authentication enabled on multiple nodes

When Tivoli Access Manager security is configured for your existing environment and security is enabled for multiple nodes, you can migrate to WebSphere Application Server, Version 6.1.

Before you begin

Note: Do not restart the WebSphere Application Server Version 7.0 servers until after performing the following procedure. The migration tools omit some files that enable the server to start correctly.

About this task

After migrating your profiles, additional steps are required when Tivoli Access Manager security is configured.

1. On the deployment manager (Host1), copy the following files from the existing directory to a comparable directory in Version 6.1:

```
%WAS_HOME%\java\jre\PDPerm.properties
%WAS_HOME%\java\jre\lib\security\PdPerm.ks
%WAS_HOME%\java\jre\PolicyDirector\PDCA.ks
%WAS_HOME%\java\jre\PolicyDirector\PD.properties
%WAS_HOME%\java\jre\PolicyDirector\PDJLog.properties
```

2. On the deployment manager, edit the PD.properties file and change the following configuration settings:

```
pd-home=C:\\Program
Files\\WebSphere\\DeploymentManager\\java\\jre\\PolicyDirector
pdvar-home=C:\\Program
Files\\WebSphere\\DeploymentManager\\java\\jre\\PolicyDirector
java-home=C:\\Program Files\\WebSphere\\DeploymentManager\\java\\jre
```

Make the appropriate changes to point to your Tivoli Access Manager Policy Server, for example:

```
pd-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
pdvar-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
java-home=C:\\Program Files\\IBM\\WebSphere\\AppServer\\java\\jre
```

3. On the deployment manager, edit the PdPerm.properties file, and change all path names to the correct path name. Change the following configuration settings:

```
pdvar-home=C:\\Program
Files\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
baseGroup.PDJv1dugong-v2dugongMessageFileHandler.fileName=C:\\Program
Files\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log\\msg__v1dugong-v2dugong.log

pdcert-url=file:/c:/progra~1/WebSphere/AppServer/java/jre/lib/security/PdPerm.ks

baseGroup.PDJv1dugong-v2dugongTraceFileHandler.fileName=C:\\Program
Files\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log\\trace__v1dugong-v2dugong.log

pd-home=C:\\Program Files\\WebSphere\\AppServer\\java\\jre\\PolicyDirector

java-home=C:\\Program Files\\WebSphere\\AppServer\\java\\jre
```

4. Start the WebSphere Application Server deployment manager.
5. On Host2, copy the following missing files from the existing directory to a comparable directory in Version 6.1:

```
%WAS_HOME%\java\jre\PDPerm.properties
%WAS_HOME%\java\jre\lib\security\PdPerm.ks
%WAS_HOME%\java\jre\PolicyDirector\PDCA.ks
```

6. On Host2, edit the PD.properties file and change the following configuration setting:

```
appsvr-plcysvrs=null\:\0:\1
```

Make the appropriate changes to point to your Tivoli Access Manager Policy Server, for example:

```
appsvr-plcysvrs=pdmgrd.test.gc.au.ibm.com\7135\1
```

7. On Host2, edit the PD.properties file, and change all path names to the correct path name. Change the following configuration settings:

```
pdvar-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
baseGroup.PDJv1dugong-v2dugongMessageFileHandler.fileName=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log\\msg__v1dugong-v2dugong.log
```

```

pdcert-url=file\:/c\:/progra~1/IBM/WebSphere/AppServer/java/jre/lib/security/PdPerm.ks

baseGroup.PDJv1dugong-v2dugongTraceFileHandler.fileName=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log/trace__v1dugong-v2dugong.log

pd-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
java-home=C:\\Program Files\\IBM\\WebSphere\\AppServer\\java\\jre
config_type=standalone

```

8. On Host2, start the node agent and its associated application server.
9. Host3, copy the following missing files from the existing directory to a comparable directory in Version 7.0:

```

%WAS_HOME%\\java\\jre\\PdPerm.properties
%WAS_HOME%\\java\\jre\\lib\\security\\PdPerm.ks
%WAS_HOME%\\java\\jre\\PolicyDirector\\PDCA.ks

```

10. On Host3, edit the Pd.properties file and change the following configuration setting:

```
appsvr-plcysvrs=null\0:1
```

Make the appropriate changes to point to your Tivoli Access Manager Policy Server, for example:

```
appsvr-plcysvrs=pdmgrd.test.gc.au.ibm.com\7135:1
```

11. On Host3, edit the PdPerm.properties file, and change all path names to the correct path name. Change the following configuration settings:

```

pdvar-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
baseGroup.PDJv1dugong-v2dugongMessageFileHandler.fileName=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log/msg__v1dugong-v2dugong.log

```

```
pdcert-url=file\:/c\:/progra~1/IBM/WebSphere/AppServer/java/jre/lib/security/PdPerm.ks
```

```

baseGroup.PDJv1dugong-v2dugongTraceFileHandler.fileName=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector\\log/trace__v1dugong-v2dugong.log

```

```

pd-home=C:\\Program
Files\\IBM\\WebSphere\\AppServer\\java\\jre\\PolicyDirector
java-home=C:\\Program Files\\IBM\\WebSphere\\AppServer\\java\\jre
config_type=standalone

```

12. On Host3, start the node agent and its associated application server.

What to do next

Also see the migration information with Tivoli Access Manager for authentication that is enabled on a single nodes with security enabled.

Preparing for security at installation time

Complete the following tasks to implement security before, during, and after installing WebSphere Application Server.

1. Secure your environment before installation. This step describes how to perform WebSphere Application Server installation with proper authority on different platforms. For more information refer to "Securing your environment before installation" on page 35.
2. Prepare the operating system for installation of WebSphere Application Server. This step describes how to prepare the different operating systems for installation of WebSphere Application Server. For more information, see "Preparing the operating system for product installation" in the InfoCenter.
3. Migrate security configurations from previous releases during installation, when you are prompted to do so. This step describes how to migrate security configurations from a previous release of WebSphere Application Server to WebSphere Application Server Version 7.0.

For more information, see "Migrating product configurations" in the InfoCenter.

4. Optional: You can create a profile during install time. If you elect to do so, administrative security is enabled for that profile "out of the box" by default. A panel is displayed during profile creation time and **enabling administrative security** is selected by default. If you elect to keep this as the default, you must supply an administrative user ID and password. This user ID is created in a federated repository, which is the default user registry when enabling administrative security at profile creation time.
5. If you go into the advanced profile creation, a panel is available for changing the default settings for your certificate, a *root certificate* (used to sign your personal certificate) and a *personal certificate* (used to sign/encrypt data over the network). Ensure that the root certificate has a long lifetime and the personal certificate a shorter one. Import your own personal certificate and or root certificate. If your personal certificate is signed by the certificate authority (CA), it is not important to change your root certificate. You should also change the default keystore password to something more secure.
6. Secure your environment after installation. This step provides information on how to protect password information after you install WebSphere Application Server. For more information, see "Securing your environment after installation."

Securing your environment before installation

The following instructions explain how to perform a product installation with proper authority.

About this task

Complete the following steps when you plan to use the local operating system as your user registry:

- **Windows** Secure your environment.
 1. Click **Start > Programs > Administrative Tools > Local Security Policy** (for domain configuration, select **Domain Security Policies**, instead).
 2. From the Local Security Settings Panel, click **Local Policies > User Rights Assignment** and add **Login as a service** rights to the login ID. The logon user must be a member of the administrator group with the rights of **Log on as a service**.
- **AIX** **HP-UX** **Linux** **Solaris** Secure your environment.
 1. Log on as **root** and verify that the umask value is **022**.
 2. To verify that the umask value is **022**, run the **umask** command.
 3. To set up the umask value as **022**, run the **umask 022** command.
 4. **Linux** **Solaris** Optional: Make sure that the /etc directory contains a shadow password file. The shadow password file is named shadow and is in the /etc directory. If the shadow password file does not exist, an error occurs after enabling administrative security and configuring the user registry as local operating system.
 5. To create the shadow file, run the **pwconv** command (without any parameters). This command creates an /etc/shadow file from the /etc/passwd file. After creating the shadow file, you can configure local operating system security.

Results

These steps help secure your environment when you use the local operating system as your user registry

Securing your environment after installation

WebSphere Application Server depends on several configuration files that are created during installation. These files contain password information and need protection. Although the files are protected to a limited degree during installation, this basic level of protection is probably not sufficient for your site. You should verify that these files are protected in compliance with the policies of your site.

Before you begin

Note: A Kerberos keytab configuration file contains a list of keys that are analogous to user passwords. The default keytab file is `krb5.keytab`. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users.

The files in the `app_server_root/profiles/profile_name/config` and `app_server_root/profiles/profile_name/properties` need protection. For example, give permission to the user who logs onto the system for WebSphere Application Server primary administrative tasks. Other users or groups, such as WebSphere Application Server console users and console groups need permissions as well.

- **Windows** Secure files on a Windows® system:
 1. Open the browser for a view of the files and directories on the machine.
 2. Locate and right-click the file or the directory that you want to protect.
 3. Click **Properties**.
 4. Click the **Security** tab.
 5. Remove the Everyone entry and any other user or group that you do not want to have access to the file.
 6. Add the users who can access the files with the proper permission.
- **UNIX** Secure files on UNIX® systems. This procedure applies only to the ordinary UNIX file system. If your site uses access-control lists, secure the files by using that mechanism. Any site-specific requirements can affect the owner, group, and corresponding privileges; for example, on the AIX® platform.
 1. Go to the `install_root` directory and change the ownership of the directory configuration and properties to the user who logs onto the system for WebSphere Application Server primary administrative tasks. Run the following command: `chown -R login_name directory_name`
Where:
 - `login_name` is a specified user or group
 - `directory_name` is the name of the directory that contains the filesIt is recommended that you assign ownership of the files that contain password information to the user who runs the application server. If more than one user runs the application server, provide permission to the group in which the users are assigned in the user registry.
 2. Set up the permission by running the following command: `chmod -R 770 directory_name`.
 3. Go to the `app_server_root/profiles/profile_name/properties` directory and set the file permissions. Set the access permissions for the following files as it pertains to your security guidelines:
 - `TraceSettings.properties`
 - `client.policy`
 - `client_types.xml`
 - `ipc.client.props`
 - `sas.client.props`
 - `sas.stdclient.properties`
 - `sas.tools.properties`
 - `soap.client.props`
 - `wsadmin.properties`
 - `wsjaas_client.conf`For example, you might issue the following command: `chmod 770 file_name` where `file_name` is the name of the file listed previously in the `install_root/profiles/profile_name/properties` directory. These files contain sensitive information such as passwords.

Note: If you enabled Kerberos authentication or SPNEGO Web authentication, set the access permissions for the following files as it pertains to your security guidelines: the Kerberos configuration file (**krb5.conf** or **krb5.ini**) and the Kerberos keytab file.

4. Create a group for WebSphere Application Server and put the users who perform full or partial WebSphere Application Server administrative tasks in that group.
5. If you want to use WebSphere MQ as a Java Messaging Service (JMS) provider, restrict access to the `/var/mqm` directories and log files used. Give write access to the user ID `mqm` or members of the `mqm` user group only.

Results

After securing your environment, only the users with permission can access the files. Failure to adequately secure these files can lead to a breach of security in your WebSphere Application Server applications.

What to do next

If failures occur that are caused by file accessing permissions, check the permission settings.

Enabling security

By enabling security, you protect your server from unauthorized users and are then able to provide application isolation and requirements for authenticating application users.

Before you begin

It is helpful to understand security from an infrastructure perspective so that you know the advantages of different authentication mechanisms, user registries, authentication protocols, and so on. Picking the right security components to meet your needs is a part of configuring security. The following sections help you make these decisions.

Read the following articles before continuing with the security configuration:

- Security

After you understand the security components, you can proceed to configure security in WebSphere Application Server.

1. Start the WebSphere Application Server administrative console.

Start the deployment manager and, in your browser, type in the address of your WebSphere Application Server Network Deployment server. By default, the console is located at `http://your_host.your_domain:9060/ibm/console`.

If security is currently disabled, you are prompted for a user ID. Log in with any user ID. However, if security is currently enabled, you are prompted for both a user ID and a password. Log in with a predefined administrative user ID and password.

2. Click **Security > Global security**. Use the Security Configuration Wizard, or configure security manually. The configuration order is not important. For more information on manual configuration, see `tsec_authusers.dita`.
3. Configure the user account repository. For more information, see “Selecting a registry or repository” on page 111. On the Global security panel, you can configure user account repositories such as federated repositories, local operating system, standalone Lightweight Directory Access Protocol (LDAP) registry, and standalone custom registry.

Note: You can choose to specify either a server ID and password for interoperability or enable a WebSphere Application Server installation to automatically generate an internal server ID. For more information about automatically generating server IDs, see “Local operating system settings” on page 118.

One of the details common to all user registries or repositories is the *Primary administrative user name*. This ID is a member of the chosen repository, but also has special privileges in WebSphere Application Server. The privileges for this ID and the privileges that are associated with the administrative role ID are the same. The Primary administrative user name can access all of the protected administrative methods.

Windows The ID must not be the same name as the machine name of your system because the repository sometimes returns machine-specific information when querying a user of the same name. In standalone LDAP registries, verify that the Primary administrative user name is a member of the repository and not just the LDAP administrative role ID. The entry must be searchable.

The Primary administrative user name does **not** run WebSphere Application Server processes. Rather, the process ID runs the WebSphere Application Server processes.

The *process ID* is determined by the way the process starts. For example, if you use a command line to start processes, the user ID that is logged into the system is the process ID. If running as a service, the user ID that is logged into the system is the user ID running the service. If you choose the local operating system registry, the process ID requires special privileges to call the operating system APIs. The process ID must have the following platform-specific privileges:

- **Windows** **Act as Part of Operating System** privileges
 - **UNIX** **Root** privileges
4. Select the **Set as current** option after you configure the user account repository. When you click **Apply** and the Enable administrative security option is set, a verification occurs to see if an administrative user ID has been configured and is present in the active user registry. The administrative user ID can be specified at the active user registry panel or from the console users link. If you do not configure an administrative ID for the active user registry, the validation fails.

Note: When you switch user registries, the `admin-authz.xml` file should be cleared of existing administrative ids and application names. Exceptions will occur in the logs for ids that exist in the `admin-authz.xml` file but do not exist in the current user registry.

5. Configure the authentication mechanism.
Configure Lightweight Third-Party Authentication (LTPA) or Kerberos, which is new to this release of WebSphere Application Server, under Authentication mechanisms and expiration. LTPA credentials can be forwarded to other machines. For security reasons, credential expire; however, you can configure the expiration dates on the console. LTPA credentials enable browsers to visit different product servers, which means you do not have to authenticate multiple times. For more information, see “Configuring the Lightweight Third Party Authentication mechanism” on page 250
If you want single sign-on (SSO) support, which provides the ability for browsers to visit different product servers without having to authenticate multiple times, see “Implementing single sign-on to minimize Web user authentications” on page 297. For form-based login, you must configure SSO when using LTPA.
6. Optional: Import and export the LTPA keys for cross-cell single Sign-on (SSO) between cells. For more information, see the following articles:
 - “Exporting Lightweight Third Party Authentication keys” on page 255.
 - “Importing Lightweight Third Party Authentication keys” on page 256
7. Configure the authentication protocol for special security requirements from Java clients, if needed. You can configure Common Secure Interoperability Version 2 (CSIV2) through links on the Global security panel. The Security Authentication Service (SAS) protocol is provided for backwards compatibility with previous product releases, but is deprecated. Links to the SAS protocol panels display on the Global security panel if your environment contains servers that use previous versions of WebSphere Application Server and support the SAS protocol. For details on configuring CSIV2 or SAS, see the article, “Configuring Common Secure Interoperability Version 2 (CSIV2) inbound and outbound communication settings” on page 466.

Note: IBM no longer ships or supports the Secure Authentication Service (SAS) IIOp security protocol. It is recommended that you use the Common Secure Interoperability version 2 (CSiv2) protocol.

8. Secure Socket Layers (SSL) is pre-configured by default, changes are not necessary unless you have custom SSL requirements. You can modify or create a new SSL configuration. This action protects the integrity of the messages sent across the Internet. The product provides a centralized location to configure SSL configurations that the various WebSphere Application Server features that use SSL can utilize, including the LDAP registry, Web container and the RMI/IIOp authentication protocol (CSiv2). For more information, see “Creating a Secure Sockets Layer configuration” on page 632. After you modify a configuration or create a new configuration, specify it on the **SSL configurations** panel. To get to the SSL configurations panel, complete the following steps:

- a. Click **Security > SSL certificate and key management**.
- b. Under Configuration settings, click **Manage endpoint security configurations > configuration_name**.
- c. Under Related items for each scope (for example, node, cluster, server), select one of the many configuration links that can be scoped to the resource you are visiting.

You can either edit the DefaultSSLConfig file or create a new SSL configuration with a new alias name. If you create a new alias name for your new keystore and truststore files, change every location that references the DefaultSSLConfig SSL configuration alias. The following list specifies the locations of where the SSL configuration repertoire aliases are used in the WebSphere Application Server configuration.

For any transports that use the new network input/output channel chains, including HTTP and Java Message Service (JMS), you can modify the SSL configuration repertoire aliases in the following locations for each server:

- Click **Server > Application server > server_name**. Under Communications, click **Ports**. Locate a transport chain where SSL is enabled and click **View associated transports**. Click *transport_channel_name*. Under Transport Channels, click **SSL Inbound Channel (SSL_2)**.
- Click **System administration > Deployment manager**. Under Additional properties, click **Ports**. Locate a transport chain where SSL is enabled and click **View associated transports**. Click *transport_channel_name*. Under Transport Channels, click **SSL Inbound Channel (SSL_2)**.
- Click **System administration > Node agents > node_agent_name**. Under Additional properties, click **Ports**. Locate a transport chain where SSL is enabled and click **View associated transports**. Click *transport_channel_name*. Under Transport Channels, click **SSL Inbound Channel (SSL_2)**.

For the Object Request Broker (ORB) SSL transports, you can modify the SSL configuration repertoire aliases in the following locations. These configurations are for the server-level for WebSphere Application Server and WebSphere Application Server Express and the cell level for WebSphere Application Server Network Deployment.

- Click **Security > Global security**. Under RMI/IIOp security, click **CSiv2 inbound communications**.
- Click **Security > Global security**. Under RMI/IIOp security, click **CSiv2 outbound communications**.

For the Lightweight Directory Access Protocol (LDAP) SSL transport, you can modify the SSL configuration repertoire aliases by clicking **Security > Global security**. Under User account repository, click the **Available realm definitions** drop-down list, and select **Standalone LDAP registry**.

9. Click **Security > Global security** to configure the rest of the security settings and enable security. For information about these settings, see “Global security settings” on page 55. For additional information, see “Server and administrative security” on page 41.
10. Validate the completed security configuration by clicking **OK** or **Apply**. If problems occur, they display at the top of the console page in red type.
11. If there are no validation problems, click **Save** to save the settings to a file that the server uses when it restarts. Saving writes the settings to the configuration repository.

Note: If you do not click **Apply** or **OK** in the Global security panel before you click **Save**, your changes are not written to the repository. The server must be restarted for any changes to take effect when you start the administrative console.

The save action enables the deployment manager to use the changed settings after WebSphere Application Server is restarted. For more information, see “Enabling security for the realm” on page 53. A Deployment manager configuration differs from a stand-alone base application server. The configuration is stored temporarily in the deployment manager until it is synchronized with all of the node agents.

Also, verify that all of the node agents are up and running in the domain. Stop all application servers during this process. If any of the node agents are down, run a manual file synchronization utility from the node agent machine to synchronize the security configuration from the deployment manager. Otherwise, the malfunctioning node agent does not communicate with the deployment manager after security is enabled on the deployment manager.

12. Start the WebSphere Application Server administrative console.

Start the deployment manager and, in your browser, type in the address of your WebSphere Application Server Network Deployment server. By default, the console is located at `http://your_host.your_domain:9060/ibm/console`.

If security is currently disabled, log in with any user ID. If security is currently enabled, log in with a predefined administrative ID and password. This ID is typically the server user ID that is specified when you configured the user registry.

Administrative security

Administrative security determines whether security is used at all, the type of registry against which authentication takes place, and other values, many of which act as defaults. Proper planning is required because incorrectly enabling administrative security can lock you out of the administrative console or cause the server to end abnormally.

Administrative security can be thought of as a “big switch” that activates a wide variety of security settings for WebSphere Application Server. Values for these settings can be specified, but they will not take effect until administrative security is activated. The settings include the authentication of users, the use of Secure Sockets Layer (SSL), and the choice of user account repository. In particular, application security, including authentication and role-based authorization, is not enforced unless administrative security is active. Administrative security is enabled by default.

Administrative security represents the security configuration that is effective for the entire security domain. A *security domain* consists of all of the servers that are configured with the same user registry *realm* name. In some cases, the realm can be the machine name of a local operating system registry. In this case, all of the application servers must reside on the same physical machine. In other cases, the realm can be the machine name of a standalone Lightweight Directory Access Protocol (LDAP) registry.

A multiple node configuration is supported because you can access remotely user registries that support the LDAP protocol. Therefore, you can enable authentication from anywhere.

The basic requirement for a security domain is that the access ID that is returned by the registry or repository from one server within the security domain is the same access ID as that returned from the registry or repository on any other server within the same security domain. The *access ID* is the unique identification of a user and is used during authorization to determine if access is permitted to the resource.

The administrative security configuration applies to every server within the security domain.

Why turn on administrative security?

Turning on administrative security activates the settings that protect your server from unauthorized users. Administrative security is enabled by default during the profile creation time. There might be some

environments where no security is needed such as a development system. On these systems you can elect to disable administrative security. However, in most environments you should keep unauthorized users from accessing the administrative console and your business applications. Administrative security must be enabled to restrict access.

What does administrative security protect?

The configuration of administrative security for a security domain involves configuring the following technologies:

- Authentication of HTTP clients
- Authentication of IIOP clients
- Administrative console security
- Naming security
- Use of SSL transports
- Role-based authorization checks of servlets, enterprise beans, and mbeans
- Propagation of identities (RunAs)
- The common user registry
- The authentication mechanism
- Other security information that defines the behavior of a security domain includes:
 - The authentication protocol (Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) security)
 - Other miscellaneous attributes

Note: It is recommended that before registering a node with an administrative agent process, that you first have administrative security enabled in the administrative agent and base profile. Once you register a profile with the administrative agent, the state of administrative security enablement cannot be changed.

Server and administrative security

The term *administrative security* refers to providing the authentication of users using the WebSphere administration functions, the use of Secure Sockets Layer (SSL), and the choice of user account repository.

In some cases, the realm can be the machine name of a Local OS user registry. In this case, all application servers must reside on the same physical machine. In other cases, the realm can be the machine name of a Lightweight Directory Access Protocol (LDAP) user registry. Because LDAP is a distributed user registry, this allows for a multiple node configuration in a Network Deployment environment. The basic requirement for a security domain is that the access ID returned by the registry from one server within the security domain is the same access ID that is returned from the registry on any other server within the same security domain. The access ID is the unique identification of a user and is used during authorization to determine if access is permitted to the resource.

Configuration of administrative security for a security domain consists of configuring the common user registry, the authentication mechanism, and other security information that defines the behavior of a security domain. The other security information that is configured includes the following components:

- Java 2 Security Manager
- Java Authentication and Authorization Service (JAAS)
- Java 2 Connector authentication data entries
- Common Secure Interoperability Version 2 (CSIv2) and Secure Authentication Service (SAS) authentication protocol (Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) security)
- Other miscellaneous attributes.

You can override some portions of the configuration at the server level.

Where multiple nodes and multiple servers within a node are possible, you can configure certain attributes at a server level. The attributes that are configurable at a server level include security enablement for the server, Java 2 security manager enablement, and CSiv2/SAS authentication protocol (RMI/IOP security). You can disable security on individual application servers while administrative security is enabled, however, you cannot enable security on an individual application server while administrative security is disabled.

While application server security is disabled for user requests, administrative and naming security is still enabled for that application server so that the administrative and naming infrastructure remains secure. If cell security is enabled, but security for individual servers is disabled, J2EE applications are not authenticated or authorized. However, naming and administrative security is still enforced. Consequently, because naming services can be called from user applications, grant Everyone access to the naming functions that are required so that these functions accept unauthenticated requests. User code does not directly access administrative security except through the supported scripting tools.

Considering security specific to a multi-node or process Network Deployment environment

WebSphere Application Server Network Deployment supports centralized management of distributed nodes and application servers. This support inherently brings complexity, especially when security is included. Because everything is distributed, security plays an even larger role in ensuring that communications are appropriately secure between application servers and node agents, and between node agents (a node-specific configuration manager) and the deployment manager (a domain-wide, centralized configuration manager).

Before you begin

Because the processes are distributed, the authentication mechanism that must be used is Lightweight Third Party Authentication (LTPA). The LTPA tokens are encrypted, signed and forwardable to remote processes. However, the tokens have expirations. The SOAP connector, which is the default connector, is used for administrative security and does not have retry logic for expired tokens. However, the protocol is stateless so a new token is created for each request if there is not sufficient time to run the request with the given time left in the token. An alternative connector is the RMI connector, which is stateful, and has some retry logic to correct expired tokens by resubmitting the requests after the error is detected. Also, because tokens have time-specific expiration, the synchronization of the system clocks is crucial to the proper operation of token-based validation. If the clocks are off by too much (approximately 10-15 minutes), you can encounter unrecoverable validation failures that can be avoided by having them in sync. Verify that the clock time, date, and time zones are all the same between systems. It is acceptable for nodes to be across time zones, provided that the times are correct within the time zones (for example, 5 PM CST = 6 PM EST, and so on).

About this task

Consider the following issues when using or planning for a Network Deployment environment.

- When attempting to run system management commands such as the **stopNode** command, explicitly specify administrative credentials to perform the operation. Most commands accept **-user** and **-password** parameters to specify the user ID and password, respectively. Specify the user ID and password of an administrative user; for example, a user who is a member of the console users with Operator or Administrator privileges or the administrative user ID that is configured in the user registry. An example of the **stopNode** command follows:

```
stopNode -username user -password pass
```
- Verify that the configuration at the node agents is always synchronized with the deployment manager prior to starting or restarting a node. To manually get the configuration synchronized, issue the

syncNode command from each node that is not synchronized. To synchronize the configuration for node agents that are started, click **System Administration > Nodes**. Select all the started nodes, and then click **Synchronize**.

- Verify that the clocks on all systems are in sync including the time zone, time and date. If they are out of sync, the tokens expire immediately when they reach the target server due to the time differences.
- Verify that the LTPA token expiration period is long enough to complete your longest downstream request. Some credentials are cached and therefore the timeout does not always include the length of the request.
- The administrative connector that is used by default for system management is SOAP. SOAP is a stateless HTTP protocol. For most situations, this connector is sufficient. If you have a problem using the SOAP connector, you might want to change the default connector on all the servers from SOAP to RMI. The RMI connector uses Common Secure Interoperability Version 2 (CSlv2), a stateful, interoperable protocol, and can be configured to use identity assertion (downstream delegation), message-layer authentication (BasicAuth or Token), and client certificate authentication (for server trust isolation). To change the default connector on a given server, go to **Administration Services** under Additional properties for that server.
- An error message might occur within the administrative subsystem security. This error indicates that the sending process did not supply a credential to the receiving process. Typically the causes of this problem are:
 - The sending process has security disabled while the receiving process has security enabled. This setup typically indicates one of the two processes are not synchronized with the cell. Having security disabled for a specific application server does not have any effect on administrative security.
 - The clocks between the systems are not synchronized; this immediately makes the credential tokens not valid. Verify that the time, date, and time zones are consistent between the two machines. An error similar to the following might occur:

```
[9/18/02 16:48:23:859 CDT] 3b9cef35 RoleBasedAuth A CWSCJ0305I: Role based authorization check failed for security name <null>, accessId NO_CRED_NO_ACCESS_ID while invoking method propagateNotifications:[Ljavax.management.Notification; on resource NotificationService and module NotificationService.
```
- When getting the following error message, validate that the clocks are synchronized between all servers within the cell, and the configurations are synchronized between all nodes and the Deployment Manager. An error similar to the following might occur:

```
[9/18/02 16:48:22:859 CDT] 3bd06f34 LTPAServerObj E CWSCJ0372E: Validation of the token failed.
```

Results

Proper understanding of the security interactions between distributed servers greatly reduces the problems that are encountered with secure communications. Security adds complexity because additional function must be managed. For security to work properly, it needs thorough consideration during the planning of your infrastructure.

What to do next

When you have security problems that are related to the WebSphere Application Server Network Deployment environment, see Chapter 14, “Troubleshooting security configurations,” on page 1289 to find additional information about the problem. When trace is needed to solve a problem because servers are distributed, it is often required to gather trace on all servers simultaneously while recreating the problem. This trace can be enabled dynamically or statically, depending on the type of problem that is occurring.

Application security

Application security enables security for the applications in your environment. This type of security provides application isolation and requirements for authenticating application users

In previous releases of WebSphere Application Server, when a user enabled global security, both administrative and application security were enabled. In WebSphere Application Server Version 6.1, the previous notion of global security is split into administrative security and application security, each of which you can enable separately.

As a result of this split, WebSphere Application Server clients must know whether application security is disabled at the target server. Administrative security is enabled, by default. Application security is disabled, by default. To enable application security, you must enable administrative security. Application security is in effect only when administrative security is enabled.

An Application Server Enablement Tag, which is specific to WebSphere Application Server, is imported into the Interoperable Object Reference (IOR) to indicate if application security is disabled for the server where the object lives. This tag is server-specific and enables clients to know when application security is disabled at the target server of its request.

For Web resources, when application security is enabled, security constraints on those resources in web.xml are enforced. When accessing a protected resource, a web client is prompted for authentication.

For enterprise bean resources, when application security is disabled, the client Common Secure Interoperability version 2 (CSIv2) code ignores the CSIv2 security tags for objects that are unknown system objects. When pure clients see that application security is disabled, these clients prompt for naming lookups, but do not prompt for enterprise bean operations.

Java 2 security

Java 2 security provides a policy-based, fine-grain access control mechanism that increases overall system integrity by checking for permissions before allowing access to certain protected system resources. Java 2 security guards access to system resources such as file I/O, sockets, and properties. Java 2 Platform, Enterprise Edition (J2EE) security guards access to Web resources such as servlets, JavaServer Pages (JSP) files and Enterprise JavaBeans™ (EJB) methods.

Because Java 2 security is relatively new, many existing or even new applications might not be prepared for the very fine-grain access control programming model that Java 2 security is capable of enforcing. Administrators need to understand the possible consequences of enabling Java 2 security if applications are not prepared for Java 2 security. Java 2 security places some new requirements on application developers and administrators.

Java 2 security for deployers and administrators

Although Java 2 security is supported, it is disabled by default. You can configure Java 2 security and administrative security independently of one another. Disabling administrative security does not disable Java 2 security automatically. You need to explicitly disable it.

If your applications, or third-party libraries are not ready, having Java 2 security enabled causes problems. You can identify these problems as Java 2 security `AccessControlExceptions` in the system log or trace files. If you are unsure about the Java 2 security readiness of your applications, disable Java 2 security initially to get your application installed and verify that it is working properly.

The policy embodied by these policy files cannot be made more restrictive because the product might not have the necessary Java 2 security `doPrivileged` APIs in place. The restrictive policy is the default policy. You can grant additional permissions, but you cannot make the default more restrictive because `AccessControlExceptions` exceptions are generated from within WebSphere Application Server. The product does not support a more restrictive policy than the default that is defined in the policy files previously mentioned.

Several policy files are used to define the security policy for the Java process. These policy files are static (code base is defined in the policy file) and in the default policy format provided by the IBM Developer Kit, Java Technology Edition. For enterprise application resources and utility libraries, WebSphere Application Server provides dynamic policy support. The code base is dynamically calculated based on deployment information and permissions are granted based on template policy files during runtime. Refer to the “Java 2 security policy files” on page 48 for more information.

Syntax errors in the policy files cause the application server process to fail, so edit these policy files carefully.

If an application is not prepared for Java 2 security, if the application provider does not provide a `was.policy` file as part of the application, or if the application provider does not communicate the expected permissions the application is likely to cause Java 2 security access control exceptions at runtime. It might not be obvious that an application is not prepared for Java 2 security. Several run-time debugging aids help troubleshoot applications that might have access control exceptions. See the Java 2 security debugging aids for more details. See “Handling applications that are not Java 2 security ready” on page 46 for information and strategies for dealing with such applications.

It is important to note when Java Security is enabled in the administrative security settings, the installed security manager does not currently check `modifyThread` and `modifyThreadGroup` permissions for non-system threads. Allowing Web and Enterprise JavaBeans (EJB) application code to create or modify a thread can have a negative impact on other components of the container and can affect the capability of the container to manage enterprise bean life cycles and transactions.

Java 2 security for application developers

Application developers must understand the permissions that are granted in the default WebSphere policy and the permission requirements of the SDK APIs that their application calls to know whether additional permissions are required. The Permissions in the Java 2 SDK reference in the resources section describes which APIs require which permission.

Application providers can assume that applications have the permissions granted in the default policy previously mentioned. Applications that access resources not covered by the default WebSphere policy are required to grant the additional Java 2 security permissions to the application.

While it is possible to grant the application additional permissions in one of the other dynamic WebSphere policy files or in one of the more traditional `java.policy` static policy files, the `was.policy` file, which is embedded in the EAR file ensures the additional permissions are scoped to the exact application that requires them. Scoping the permission beyond the application code that requires it can permit code that normally does not have permission to access particular resources.

If an application component is being developed, like a library that might actually be included in more than one `.ear` file, then the library developer needs to document the required Java 2 permissions that are required by the application assembler. There is no `was.policy` file for library-type components. The developer must communicate the required permissions through application programming interface (API) documentation or some other external documentation.

If the component library is shared by multiple enterprise applications, the permissions can be granted to all enterprise applications on the node in the `app.policy` file.

Note: Updates to the `app.policy` file only apply to the enterprise applications on the node to which the `app.policy` file belongs.

If the permission is only used internally by the component library and the application is never granted access to resources that are protected by the permission, it might be necessary to mark the code as

privileged. Refer to the, `AccessControlException`, topic for more details. However, improperly inserting a `doPrivileged` call might open up security holes. Understand the implication of `doPrivileged` call to make a correct judgement.

The section on Dynamic policy files in “Java 2 security policy files” on page 48 describes how the permissions in the `was.policy` files are granted at runtime.

Debugging Aids

The WebSphere Application Server `SYSOUT` file and the `com.ibm.websphere.java2secman.norethrow` property are the two primary aids for debugging.

The WebSphere System Log or Trace Files

The `AccessControl` exception that is logged in the system log or trace files contains the permission violation that causes the exception, the exception call stack, and the permissions granted to each stack frame. This information is usually enough to determine the missing permission and the code requiring the permission.

The `com.ibm.websphere.java2secman.norethrow` property

When Java 2 security is enabled in WebSphere Application Server, the security manager component creates a `java.security.AccessControl` exception when a permission violation occurs. This exception, if not handled, often causes a run-time failure. This exception is also logged in the `SYSOUT` file.

However, when the Java virtual machine `com.ibm.websphere.java2secman.norethrow` property is set and has a value of `true`, the security manager does not create the `AccessControl` exception. This information is logged.

This property is intended for a sandbox or debug environment because it instructs the security manager not to create the `AccessControl` exception. Java 2 security is not enforced. Do not use this property in a production environment where a relaxed Java 2 security environment weakens the integrity that Java 2 security is intended to produce.

This property is valuable in a sandbox or test environment where the application can be thoroughly tested and where the system log or trace files can be inspected for `AccessControl` exceptions. Because this property does not create the `AccessControl` exception, it does not propagate the call stack and does not cause a failure. Without this property, you have to find and fix `AccessControl` exceptions one at a time.

Handling applications that are not Java 2 security ready

If the increased system integrity that Java 2 security provides is important, then contact the application provider to have the application support Java 2 security or at least communicate the required additional permissions beyond the default WebSphere Application Server policy that must be granted.

The easiest way to deal with such applications is to disable Java 2 security in WebSphere Application Server. The downside is that this solution applies to the entire system and the integrity of the system is not as strong as it might be. Disabling Java 2 security might not be acceptable depending on the organization security policies or risk tolerances.

Another approach is to leave Java 2 security enabled, but to grant either just enough additional permissions or grant all permissions to just the problematic application. Granting permissions however, might not be a trivial thing to do. If the application provider has not communicated the required permissions in some way, no easy way exists to determine what the required permissions are and granting all permissions might be the only choice. You minimize this risk by locating this application on a different

node, which might help isolate it from certain resources. Grant the `java.security.AllPermission` permission in the `was.policy` file that is embedded in the application `.ear` file, for example:

```
grant codeBase "file:${application}" {
    permission java.security.AllPermission;
};
```

The server.policy file

The `server.policy` file is located in the `app_server_root/properties/` directory.

This policy defines the policy for the WebSphere Application Server classes. At present, all the server processes on the same installation share the same `server.policy` file. However, you can configure this file so that each server process can have a separate `server.policy` file. Define the policy file as the value of the `java.security.policy` Java system properties. For details of how to define Java system properties, refer to the Process definition section of the Manage application servers file.

The `server.policy` file is not a configuration file managed by the repository and the file replication service. Changes to this file are local and do not get replicated to other machines. Use the `server.policy` file to define Java 2 security policy for server resources. Use the `app.policy` file (per node) or the `was.policy` file (per enterprise application) to define Java 2 security policy for enterprise application resources.

Note: Updates to the `app.policy` file only apply to the enterprise applications on the node to which the `app.policy` file belongs.

The java.policy file

The file represents the default permissions that are granted to all classes. The policy of this file applies to all the processes launched by the Java Virtual Machine in the WebSphere Application Server.

The `java.policy` file is located in the `app_server_root/java/lib/security` directory.

Troubleshooting

Symptom:

Error message CWSCJ0314E: Current Java 2 security policy reported a potential violation of Java 2 security permission. Refer to Problem Determination Guide for further information. {0}Permission\:{1}Code\:{2}{3}Stack Trace\:{4}Code Base Location\:{5} Current® Java 2 security policy reported a potential violation of Java 2 Security Permission. Refer to Problem Determination Guide for further information. {0}Permission\:{1}Code\:{2}{3}Stack Trace\:{4}Code Base Location\:{5}

Problem:

The Java security manager `checkPermission` method reported a security exception on the subject permission with debugging information. The reported information can be different with respect to the system configuration. This report is enabled by either configuring a Reliability Availability Service Ability (RAS) trace into debug mode or specifying a Java property.

See Enabling trace for information on how to configure RAS trace in debug mode.

Specify the following property in the JVM Settings panel from the administrative console:

java.security.debug. Valid values include:

access

Print all debug information including: required permission, code, stack, and code base location.

stack Print debug information including: required permission, code, and stack.

failure Print debug information including: required permission and code.

Recommended response:

The reported exception might be critical to the secure system. Turn on security trace to determine the potential code that might have violated the security policy. After the violating code is determined, verify if the attempted operation is permitted with respect to Java 2 security, by examining all applicable Java 2 security policy files and the application code.

If the application is running with Java Mail, this message might be benign. You can update the `was.policy` file to grant the following permissions to the application:

```
permission java.io.FilePermission "${user.home}${}/.mailcap", "read";
permission java.io.FilePermission "${user.home}${}/.mime.types", "read";
permission java.io.FilePermission "${java.home}${}/lib${}/mailcap", "read";
permission java.io.FilePermission "${java.home}${}/lib${}/mime.types", "read";
```

Messages

Message:	CWSCJ0313E: Java 2 security manager debug message flags are initialized\: TrDebug: {0}, Access: {1}, Stack: {2}, Failure: {3}
Problem:	Configured values of the valid debug message flags for security manager.

Message:	CWSCJ0307E: Unexpected exception is caught when trying to determine the code base location. Exception: {0}
Problem:	An unexpected exception is caught when the code base location is determined.

Java 2 security policy files

The Java 2 Platform, Enterprise Edition (J2EE) Version 1.3 and later specifications have a well-defined programming model of responsibilities between the container providers and the application code. Using Java 2 security manager to help enforce this programming model is recommended. Certain operations are not supported in the application code because such operations interfere with the behavior and operation of the containers. The Java 2 security manager is used in the product to enforce responsibilities of the container and the application code.

This product provides support for policy file management. A number of policy files in the product are either static or dynamic. *Dynamic policy* is a template of permissions for a particular type of resource. No relative code base is defined in the dynamic policy template. The code base is dynamically calculated from the deployment and run-time data.

Static policy files

Policy file	Location
java.policy	<i>app_server_root</i> /java/jre/lib/security/java.policy. Default permissions are granted to all classes. The policy of this file applies to all the processes launched by WebSphere Application Server.
server.policy	<i>profile_root</i> /properties/server.policy. Default permissions are granted to all the product servers.
client.policy	<i>profile_root</i> /properties/client.policy. Default permissions are granted for all of the product client containers and applets on a node.

The static policy files are not managed by configuration and file replication services. Changes made in these files are local and are not replicated to other nodes in the Network Deployment cell.

Dynamic policy files

Policy file	Location
spi.policy	<p><i>profile_root/config/cells/cell_name</i> <i>/nodes/node_name/spi.policy</i></p> <p>This template is for the Service Provider Interface (SPI) or the third-party resources that are embedded in the product. Examples of SPI are the Java Message Service (JMS) in MQ Series and Java database connectivity (JDBC) drivers. The code base for the embedded resources are dynamically determined from the configuration (<i>resources.xml</i> file) and run-time data, and permissions that are defined in the <i>spi.policy</i> files are automatically applied to these resources and JAR files that are specified in the class path of a resource adapter. The default permission of the <i>spi.policy</i> file is <code>java.security.AllPermissions</code>.</p>
library.policy	<p><i>profile_root/config/cells/cell_name/nodes</i> <i>/node_name/library.policy</i></p> <p>This template is for the library (Java library classes). You can define a shared library to use in multiple product applications. The default permission of the <i>library.policy</i> file is empty.</p>
app.policy	<p><i>profile_root/config/cells/cell_name</i> <i>/nodes/node_name/app.policy</i></p> <p>The <i>app.policy</i> file defines the default permissions that are granted to all of the enterprise applications running on <i>node_name</i> in <i>cell_name</i>. Note: Updates to the <i>app.policy</i> file only apply to the enterprise applications on the node to which the <i>app.policy</i> file belongs.</p>
was.policy	<p><i>profile_root/config/cells/cell_name</i> <i>/applications/ear_file_name/deployments/</i> <i>application_name/META-INF/was.policy</i></p> <p>This template is for application-specific permissions. The <i>was.policy</i> file is embedded in the enterprise archive (EAR) file.</p>
ra.xml	<p><i>rar_file_name/META-INF/was.policy.RAR.</i></p> <p>This file can have a permission specification that is defined in the <i>ra.xml</i> file. The <i>ra.xml</i> file is embedded in the RAR file.</p>

Grant entries that are specified in the *app.policy* and *was.policy* files must have a code base defined. If grant entries are specified without a code base, the policy files are not loaded properly and the application can fail. If the intent is to grant the permissions to all applications, use `file:${application}` as a code base in the grant entry.

Syntax of the policy file

A policy file contains several policy entries. The following example depicts each policy entry format:

```
grant [codebase <Codebase>] {
  permission <Permission>;
  permission <Permission>;
  permission <Permission>;
};
```

<CodeBase>: A URL.

For example, `"file:${java.home}/lib/tools.jar"`

When [codebase <Codebase>] is not specified, listed permissions are applied to everything.

If URL ends with a JAR file name, only the classes in the JAR file belong to the codebase.

If URL ends with "/", only the class files in the specified directory belong to the codebase.

If URL ends with "*", all JAR and class files in the specified directory belong to the codebase.
 If URL ends with "-", all JAR and class files in the specified directory and its subdirectories belong to the codebase.

<Permissions>: Consists from
 Permission Type : class name of the permission
 Target Name : name specifying the target
 Actions : actions allowed on target
 For example,
 java.io.FilePermission "/tmp/xxx", "read,write"

Refer to developer kit specifications for the details of each permission.

Syntax of dynamic policy

You can define permissions for specific types of resources in dynamic policy files for an enterprise application. This action is achieved by using *product-reserved symbols*. The reserved symbol scope depends on where it is defined. If you define the permissions in the app.policy file, the symbol applies to all the resources on all of the enterprise applications that run on *node_name*. If you define the permissions in the META-INF/was.policy file, the symbol applies only to the specific enterprise application. Valid symbols for the code base are listed in the following table:

Symbol	Meaning
file:\${application}	Permissions apply to all the resources within the application
file:\${jars}	Permissions apply to all the utility Java archive (JAR) files within the application
file:\${ejbComponent}	Permissions apply to the Enterprise JavaBeans (EJB) resources within the application
file:\${webComponent}	Permissions apply to the Web resources within the application
file:\${connectorComponent}	Permissions apply to the connector resources within the application

You can specify the module name for a granular setting, except for these entries that are specified by the code base symbols. For example:

```
grant codeBase "file:DefaultWebApplication.war" {
    permission java.security.SecurityPermission "printIdentity";
};

grant codeBase "file:IncCMP11.jar" {
    permission java.io.FilePermission
    "${user.install.root}${/}bin${/}DefaultDB${/}-",
    "read,write,delete";
};
```

The sixth and seventh lines in the previous code sample are one continuous line. You can use a relative code base only in the META-INF/was.policy file. Several product-reserved symbols are defined to associate the permission lists to a specific type of resources.

Symbol	Meaning
file:\${application}	Permissions apply to all the resources within the application
file:\${jars}	Permissions apply to all the utility JAR files within the application

Symbol	Meaning
file:\${ejbComponent}	Permissions apply to the enterprise beans resources within the application
file:\${webComponent}	Permissions apply to the Web resources within the application
file:\${connectorComponent}	Permissions apply to the connector resources both within the application and in the standalone connector resources.

Five embedded symbols are provided to specify the path and the name for the `java.io.FilePermission` permission. These symbols enable flexible permission specification. The absolute file path is fixed after the installation of the application.

Symbol	Meaning
\${app.installed.path}	Path where the application is installed
\${was.module.path}	Path where the module is installed
\${current.cell.name}	Current cell name
\${current.node.name}	Current node name
\${current.server.name}	Current server name

Note: Do not use the `${was.module.path}` in the `${application}` entry.

Carefully determine where to add a new permission. An incorrectly specified permission causes an `AccessControlException` exception. Because dynamic policy resolves the code base at runtime, determining which policy file has a problem is difficult. Add a permission only to the necessary resources. For example, use `${ejbcomponent}`, and etc instead of `${application}`, and update the `was.policy` file instead of the `app.policy` file, if possible.

Static policy filtering

Limited static policy filtering support exists. If the `app.policy` file and the `was.policy` file have permissions that are defined in the `filter.policy` file with the `filterMask` keyword, the runtime removes the permissions from the applications and an audit message is logged. However, if the permissions that are defined in the `app.policy` and the `was.policy` files are compound permissions, for example, `java.security.AllPermission`, the permission is not removed, but a warning message is written to the log file. The policy filtering only supports Developer Kit permissions; the permissions package name begins with `java` or `javax`.

Run-time policy filtering support is provided to force stricter filtering. If the `app.policy` file and the `was.policy` file have permissions that are defined in the `filter.policy` file with the `runtimeFilterMask` keyword, the runtime removes the permissions from the applications no matter what permissions are granted to the application. For example, even if a `was.policy` file has the `java.security.AllPermission` permission granted to one of its modules, specified permissions such as the `runtimeFilterMask` permission are removed from the granted permission during runtime.

Policy file editing

Using the policy tool that is provided by the Developer Kit (`app_server_root/java/jre/bin/policytool`), to edit the previous policy files is recommended. For Network Deployment, extract the policy files from the repository before editing. After the policy file is extracted, use the policy tool to edit the file. Check the modified policy files into the repository and synchronize them with other nodes.

Troubleshooting

To debug the dynamic policy, choose one of three ways to generate the detail report of the `AccessControlException` exception.

- **Trace** (Configured by RAS trace). Enables traces with the trace specification:

Note: The following command is one continuous line

```
com.ibm.ws.security.policy.*=all=enabled:  
com.ibm.ws.security.core.SecurityManager=all=enabled
```

- **Trace** (Configured by property). Specifies a Java `java.security.debug` property. Valid values for the `java.security.debug` property are as follows:
 - Access. Print all debug information including required permission, code, stack, and code base location.
 - Stack. Print debug information including, required permission, code, and stack.
 - Failure. Print debug information including required permission and code.
- **ffdc**. Enable `ffdc`, modify the `ffdcRun.properties` file by changing `Level=4` and `LAE=true`. Look for an Access Violation keyword in the log file.

Access control exception

The Java 2 security behavior is specified by its *security policy*. The security policy is an access-control matrix that specifies which system resources certain code bases can access and who must sign them. The Java 2 security policy is declarative and it is enforced by the `java.security.AccessController.checkPermission` method.

The following example depicts the algorithm for the `java.security.AccessController.checkPermission` method. For the complete algorithm, refer to the Java 2 security check permission algorithm in Resources for learning.

```
i = m;  
while (i > 0) {  
    if (caller i's domain does not have the permission)  
        throw AccessControlException;  
    else if (caller i is marked as privileged)  
        return;  
    i = i - 1;  
};
```

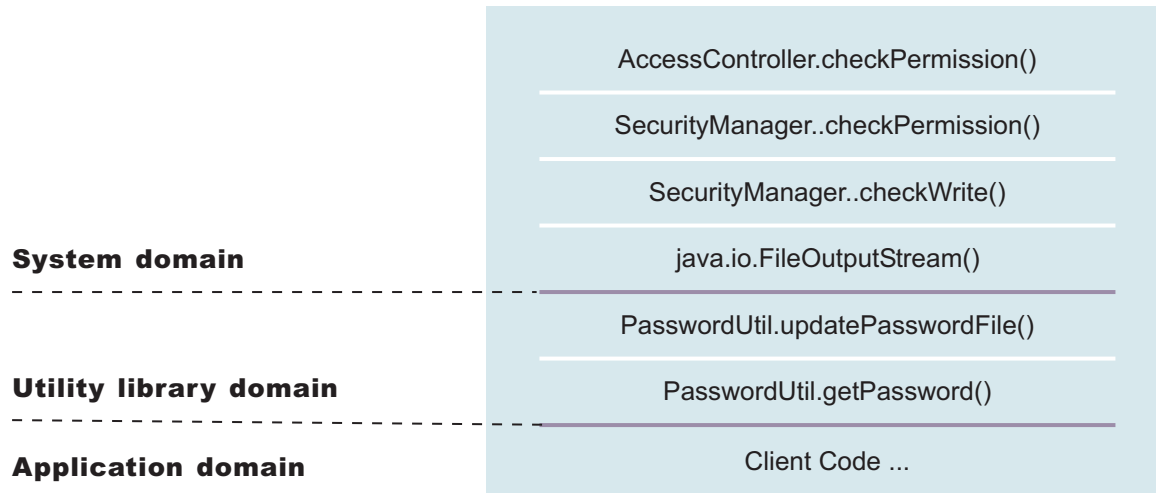
The algorithm requires that all the classes or callers on the call stack have the permissions when a `java.security.AccessController.checkPermission` method is performed or the request is denied and a `java.security.AccessControlException` exception is created. However, if the caller is marked as *privileged* and the class (caller) is granted these permissions, the algorithm returns and does not traverse the entire call stack. Subsequent classes (callers) do not need the required permission granted.

A `java.security.AccessControlException` exception is created when certain classes on the call stack are missing the required permissions during a `java.security.AccessController.checkPermission` method. Two possible resolutions to the `java.security.AccessControlException` exception are as follows:

- If the application is calling a Java 2 security-protected application programming interface (API), grant the required permission to the application Java 2 security policy. If the application is not calling a Java 2 security-protected API directly, the required permission results from the side-effect of the third-party APIs accessing Java 2 security-protected resources.
- If the application is granted the required permission, it gains more access than it needs. In this case, it is likely that the third party code that accesses the Java 2 security-protected resource is not properly marked as *privileged*.

Example call stack

This example of a call stack indicates where application code is using a third-party API utility library to update the password. The following example is presented to illustrate the point. The decision of where to mark the code as privileged is application-specific and is unique in every situation. This decision requires great depth of domain knowledge and security expertise to make the correct judgement. A number of well written publications and books are available on this topic. Referencing these materials for more detailed information is recommended.



You can use the **PasswordUtil** utility to change the password of a user. The utility types in the old password and the new password twice to ensure that the correct password is entered. If the old password matches the one stored in the password file, the new password is stored and the password file updates. Assume that none of the stack frame is marked as privileged. According to the `java.security.AccessController.checkPermission` algorithm, the application fails unless all the classes on the call stack are granted write permission to the password file. The client application does not have permission to write to the password file directly and to update the password file at will.

However, if the `PasswordUtil.updatePasswordFile` method marks the code that accesses the password file as privileged, then the check permission algorithm does not check for the required permission from classes that call the `PasswordUtil.updatePasswordFile` method for the required permission as long as the **PasswordUtil** class is granted the permission. The client application can successfully update a password without granting the permission to write to the password file.

The ability to mark code privileged is very flexible and powerful. If this ability is used incorrectly, the overall security of the system can be compromised and security holes can be exposed. Use the ability to mark code privileged carefully.

Resolution to the `java.security.AccessControlException` exception

As described previously, you have two approaches to resolve a `java.security.AccessControlException` exception. Judge these exceptions individually to decide which of the following resolutions is best:

1. Grant the missing permission to the application.
2. Mark some code as privileged, after considering the issues and risks.

Enabling security for the realm

Use this topic to enable IBM WebSphere Application Server security. You must enable administrative security for all other security settings to function.

About this task

WebSphere Application Server uses cryptography to protect sensitive data and to ensure confidentiality and integrity of communications between WebSphere Application Server and other components in the network. Cryptography is also used by Web services security when certain security constraints are configured for the Web Services application.

WebSphere Application Server uses Java Secure Sockets Extension (JSSE) and Java Cryptography Extension (JCE) libraries in the Software Development Kit (SDK) to perform this cryptography. The SDK provides strong but limited jurisdiction policy files. Unrestricted policy files provide the ability to perform full strength cryptography and to improve performance.

WebSphere Application Server provides a SDK 6 that contains strong, but limited jurisdiction policy files. You can download the unrestricted policy files from the following Web site: IBM developer kit: Security information. Complete the following steps to download and install the new policy files:

1. Click **Java SE 6**
2. Scroll down the page then click **IBM SDK Policy files**.
The Unrestricted JCE Policy files for SDK 6 Web site displays.
3. Click **Sign in** and provide your IBM.com ID and password.
4. Select **Unrestricted JCE Policy files for SDK 6** and click **Continue**.
5. View the license and click **I Agree** to continue.
6. Click **Download Now**.
7. Extract the unlimited jurisdiction policy files that are packaged in the ZIP file. The ZIP file contains a `US_export_policy.jar` file and a `local_policy.jar` file.
8. In your WebSphere Application Server installation, go to the `$JAVA_HOME/jre/lib/security` directory and back up your `US_export_policy.jar` and `local_policy.jar` files.
9. Replace your `US_export_policy.jar` and `local_policy.jar` files with the two files that you downloaded from the IBM.com Web site.
1. Enable administrative security in WebSphere Application Server.
For more information, see “Enabling security” on page 37. It is important to click **Security > Global security**. Select an available realm definition from the list, and then click **Set as current**. Save the configuration to the repository. Verify that the validation that occurs after you click **OK** on the **Security > Global security** panel is successful before continuing. If the validation is not successful and you continue with these steps, you risk the server not starting. Re-configure the security settings until validation is successful.
2. Send a copy of the new configuration to all of the running node agents using the administrative console. If a node agent fails to get the security-enabled configuration, communication with the deployment manager fails, due to a lack of access. The node agent is not security-enabled. To force synchronize a specific node, complete the following steps from the administrative console:
 - a. Click **System administration > Nodes** and select the option next to all the nodes. You do not need to select the deployment manager node.
 - b. Click **Full resynchronize** to verify that the file synchronization has occurred. The message might indicate that the nodes already are synchronized. This message is OK. When synchronization is initiated, verify that the Synchronized status displays for all nodes.
3. Stop the deployment manager. Manually restart the deployment manager from the command line or service. To stop the deployment manager, click **System administration > Deployment manager** and click **Stop**. This action logs you out of the administrative console and stops the deployment manager process.
4. Restart the deployment manager process.
To restart the deployment manager process, locate the `app_server_root/bin` directory and type `startManager.bat` or `startManager.sh`. After the deployment manager initialization is complete, go back

into the administrative console to complete this task. Remember that security now is enabled in only the deployment manager. If you enabled single sign-on (SSO), specify the fully qualified domain name of your Web address, for example, `http://myhost.domain:port_number/ibm/console`. When you are prompted for a user ID and password, type the one that you defined as the administrator ID in the configured user registry.

5. If the deployment manager does not start after enabling security, disable security using a script and restart. Disable security by issuing the following command from the `DeploymentManager/bin` directory:

```
./wsadmin.sh -conntype NONE
```

At the prompt, enter `securityoff`.

6. Restart all node agents to make them security enabled. You must have restarted the deployment manager in a previous step before completing this step. If the node agent is security-enabled before the deployment manager is security-enabled, the deployment manager cannot query the node agent for status or give the node agent commands. To stop all node agents, complete the following steps:
 - a. Go to **System administration > Node agents** and select the option beside all node agents. Click **Restart**. A message similar to the following example is displayed at the top of the panel: The node agent on node NODE NAME was restarted successfully.
 - b. Alternatively, if you previously did not stop your application servers, restart all of the servers within any given node by clicking **System administration > Node agents** and by clicking the node agents where you want to restart all the servers. Click **Restart all Servers on Node**. This action restarts the node agent and any started application servers.
7. If any node agent fails to restart, perform a manual resynchronization of the configuration. This step consists of going to the physical node and running the client **syncNode** command. This client logs into the deployment manager and copies all of the configuration files to the node agent. This action ensures that the configuration is security-enabled. If the node agent is started, but is not communicating with the deployment manager, stop the node agent by issuing the **stopServer** command.

Global security settings

Use this panel to configure administration and the default application security policy. This security configuration applies to the security policy for all administrative functions and is used as a default security policy for user applications. Security domains can be defined to override and customize the security policies for user applications.

To view this administrative console page, click **Security > Global security**.

Security has some performance impacts on your applications. The performance impacts can vary depending upon the application workload characteristics. You must first determine that the needed level of security is enabled for your applications, and then measure the impact of security on the performance of your applications.

When security is configured, validate any changes to the user registry or authentication mechanism panels. Click **Apply** to validate the user registry settings. An attempt is made to authenticate the server ID or to validate the admin ID (if `internalServerID` is used) to the configured user registry. Validating the user registry settings after enabling administrative security can avoid problems when you restart the server for the first time.

Security configuration wizard:

Launches a wizard that enables you to configure the basic administrative and application security settings. This process restricts administrative tasks and applications to authorized users.

Using this wizard, you can configure application security, resource or Java 2 Connector (J2C) security, and a user registry. You can configure an existing registry and enable administrative, application, and resource security.

When you apply changes made by using the security configuration wizard, administrative security is turned on by default.

Security configuration report:

Launches a security configuration report that displays the core security settings of the application server. The report also displays the administrative users and groups and the CORBA naming roles.

A current limitation to the report is that it does not display application level security information. The report also does not display information on Java Message Service (JMS) security, bus security, or Web Services security.

When multiple security domains are configured, the report displays the security configuration associated with each domain.

Enable administrative security:

Specifies whether to enable administrative security for this application server domain. Administrative security requires users to authenticate before obtaining administrative control of the application server.

For more information, see the related links for administrative roles and administrative authentication.

When enabling security, set the authentication mechanism configuration and specify a valid user ID and password (or a valid admin ID when internalServerID feature is used) in the selected registry configuration.

Note: There is a difference between the user ID (which is normally called the admin ID), which identifies administrators who manage the environment, and a server ID, which is used for server-to-server communication. You do not need to enter a server ID and password when you are using the internal server ID feature. However, optionally, you can specify a server ID and password. To specify the server ID and password, complete the following steps:

1. Click **Security > Global security**.
2. Under User accounts repository, select the repository and click **Configure**.
3. Specify the server ID and password in the Server user identity section.

If you have problems, such as the server not starting after enabling security within the security domain, resynchronize all of the files from the cell to this node. To resynchronize files, run the following command from the node: `syncNode -username your_userid -password your_password`. This command connects to the deployment manager and resynchronizes all of the files.

Default: Enabled

Enable application security:

Enables security for the applications in your environment. This type of security provides application isolation and requirements for authenticating application users

In previous releases of WebSphere Application Server, when a user enabled global security, both administrative and application security were enabled. In WebSphere Application Server Version 6.1, the previous notion of global security is split into administrative security and application security, each of which you can enable separately.

As a result of this split, WebSphere Application Server clients must know whether application security is disabled at the target server. Administrative security is enabled, by default. Application security is disabled, by default. To enable application security, you must enable administrative security. Application security is in effect only when administrative security is enabled.

Default: Disabled

Use Java 2 security to restrict application access to local resources:

Specifies whether to enable or disable Java 2 security permission checking. By default, access to local resources is not restricted. You can choose to disable Java 2 security, even when application security is enabled.

When the **Use Java 2 security to restrict application access to local resources** option is enabled and if an application requires more Java 2 security permissions than are granted in the default policy, the application might fail to run properly until the required permissions are granted in either the app.policy file or the was.policy file of the application. AccessControl exceptions are generated by applications that do not have all the required permissions. See the related links for more information about Java 2 security.

Default: Disabled

Warn if applications are granted custom permissions:

Specifies that during application deployment and application start, the security runtime issues a warning if applications are granted any custom permissions. Custom permissions are permissions that are defined by the user applications, not Java API permissions. Java API permissions are permissions in the java.* and javax.* packages.

The application server provides support for policy file management. A number of policy files are available in this product, some of them are static and some of them are dynamic. Dynamic policy is a template of permissions for a particular type of resource. No code base is defined and no relative code base is used in the dynamic policy template. The real code base is dynamically created from the configuration and run-time data. The filter.policy file contains a list of permissions that you do not want an application to have according to the J2EE 1.4 specification. For more information on permissions, see the related link about Java 2 security policy files.

Note: You cannot enable this option without enabling the **Use Java 2 security to restrict application access to local resources** option.

Default: Disabled

Restrict access to resource authentication data:

Enable this option to restrict application access to sensitive Java Connector Architecture (JCA) mapping authentication data.

Consider enabling this option when both of the following conditions are true:

- Java 2 security is enforced.
- The application code is granted the accessRuntimeClasses WebSphereRuntimePermission permission in the was.policy file found within the application enterprise archive (EAR) file. For example, the application code is granted the permission when the following line is found in your was.policy file:
permission com.ibm.websphere.security.WebSphereRuntimePermission "accessRuntimeClasses";

The **Restrict access to resource authentication data** option adds fine-grained Java 2 security permission checking to the default principal mapping of the WSPrincipalMappingLoginModule implementation. You must grant explicit permission to Java 2 Platform, Enterprise Edition (J2EE) applications that use the WSPrincipalMappingLoginModule implementation directly in the Java Authentication and Authorization Service (JAAS) login when **Use Java 2 security to restrict application**

access to local resources and the **Restrict access to resource authentication data** options are enabled.

Default: Disabled

Current realm definition:

Specifies the current setting for the active user repository.

This field is read-only.

Available realm definitions:

Specifies the available user account repositories.

The selections appear in a drop-down list containing:

- Local operating system
- Standalone LDAP registry
- Standalone custom registry

Set as current:

Enables the user repository after it is configured.

You can configure settings for one of the following user repositories:

Federated repositories

Specify this setting to manage profiles in multiple repositories under a single realm. The realm can consist of identities in:

- The file-based repository that is built into the system
- One or more external repositories
- Both the built-in, file-based repository and in one or more external repositories

Note: Only a user with administrator privileges can view the federated repositories configuration.

Local operating system

You cannot use localOS in multi-node or when running as non-root on a UNIX platform.

The local operating system registry is valid only when you use a domain controller or the Network Deployment cell resides on a single machine. In the later case, you cannot spread multiple nodes in a cell across multiple machines as this configuration, using the local OS user registry, is not valid.

Standalone LDAP registry

Specify this setting to use standalone LDAP registry settings when users and groups reside in an external LDAP directory. When security is enabled and any of these properties change, go to the **Security > Global security** panel and click **Apply** or **OK** to validate the changes.

Note: Since multiple LDAP servers are supported, this setting does not imply one LDAP registry.

Standalone custom registry

Specify this setting to implement your own standalone custom registry that implements the `com.ibm.websphere.security.UserRegistry` interface. When security is enabled and any of these properties change, go to the Global security panel and click **Apply** or **OK** to validate the changes.

Default: Disabled

Configure...:

Select to configure the global security settings.

Web and SIP security:

Under Authentication, expand Web and SIP security to view links to:

- General settings
- Single sign-on (SSO)
- SPNEGO web authentication
- Trust association

General settings:

Select to specify the settings for Web authentication.

Single sign-on (SSO):

Select to specify the configuration values for single sign-on (SSO).

With SSO support, Web users can authenticate once when accessing both WebSphere Application Server resources, such as HTML, JavaServer Pages (JSP) files, servlets, enterprise beans, and Lotus® Domino® resources.

SPNEGO web authentication:

Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) provides a way for Web clients and the server to negotiate the web authentication protocol that is used to permit communications.

Trust association:

Select to specify the settings for the trust association. Trust association is used to connect reversed proxy servers to the application servers.

You can use the global security settings or customize the settings for a domain.

Note: The use of trust association interceptors (TAIs) for SPNEGO authentication is now deprecated. The SPNEGO web authentication panels now provide a much easier way to configure SPNEGO.

RMI/IIOP security:

Under Authentication, expand RMI/IIOP security to view links to:

- CSiv2 inbound communications
- CSiv2 outbound communications

CSiv2 inbound communications:

Select to specify authentication settings for requests that are received and transport settings for connections that are accepted by this server using the Object Management Group (OMG) Common Secure Interoperability (CSI) authentication protocol.

Authentication features include three layers of authentication that you can use simultaneously:

- **CSlv2 attribute layer.** The attribute layer might contain an identity token, which is an identity from an upstream server that already is authenticated. The identity layer has the highest priority, followed by the message layer, and then the transport layer. If a client sends all three, only the identity layer is used. The only way to use the SSL client certificate as the identity is if it is the only information that is presented during the request. The client picks up the interoperable object reference (IOR) from the namespace and reads the values from the tagged component to determine what the server needs for security.
- **CSlv2 transport layer.** The transport layer, which is the lowest layer, might contain a Secure Sockets Layer (SSL) client certificate as the identity.
- **CSlv2 message layer.** The message layer might contain a user ID and password or an authenticated token with an expiration.

CSlv2 outbound communications:

Select to specify authentication settings for requests that are sent and transport settings for connections that are initiated by the server using the Object Management Group (OMG) Common Secure Interoperability (CSI) authentication protocol.

Authentication features include three layers of authentication that you can use simultaneously:

- **CSlv2 attribute layer.** The attribute layer might contain an identity token, which is an identity from an upstream server that already is authenticated. The identity layer has the highest priority, followed by the message layer, and then the transport layer. If a client sends all three, only the identity layer is used. The only way to use the SSL client certificate as the identity is if it is the only information that is presented during the request. The client picks up the interoperable object reference (IOR) from the namespace and reads the values from the tagged component to determine what the server needs for security.
- **CSlv2 transport layer.** The transport layer, which is the lowest layer, might contain a Secure Sockets Layer (SSL) client certificate as the identity.
- **CSlv2 message layer.** The message layer might contain a user ID and password or an authenticated token with an expiration.

Java authentication and authorization service:

Under Authentication, expand Java authentication and authorization service to view links to:

- Application logins
- System logins
- J2C authentication data

Application logins:

Select to define login configurations that are used by JAAS.

Do not remove the ClientContainer, DefaultPrincipalMapping, and WSLogin login configurations because other applications might use them. If these configurations are removed, other applications might fail.

System logins:

Select to define the JAAS login configurations that are used by system resources, including the authentication mechanism, principal mapping, and credential mapping.

J2C authentication data:

Select to specify the settings for the Java Authentication and Authorization Service (JAAS) Java 2 Connector (J2C) authentication data.

You can use the global security settings or customize the settings for a domain.

LTPA:

Select to encrypt authentication information so that the application server can send the data from one server to another in a secure manner.

The encryption of authentication information that is exchanged between servers involves the Lightweight Third-Party Authentication (LTPA) mechanism.

Kerberos and LTPA:

Select to encrypt authentication information so that the application server can send the data from one server to another in a secure manner.

The encryption of authentication information that is exchanged between servers involves the Kerberos mechanism.

Note: Kerberos must be configured before this option can be selected.

Kerberos configuration:

Select to encrypt authentication information so that the application server can send data from one server to another in a secure manner.

The encryption of the authentication information that is exchanged between servers involves the KRB5 of LTPA mechanism.

Authentication cache settings:

Select to set your authentication cache settings.

Use realm-qualified user names:

Specifies that user names that are returned by methods, such as the `getUserPrincipal()` method, are qualified with the security realm in which they reside.

Security domains:

Use the Security Domain link to configure additional security configurations for user applications.

For example, if you want use a different user registry for a set of user applications than the one used at the global level, you can create a security configuration with that user registry and associate it with that set of applications. These additional security configurations can be associated with various scopes (cell, clusters/servers, SIBuses). Once the security configurations have been associated with a scope all of the user applications in that scope use this security configuration. Read about “Multiple security domains” on page 78 for more detailed information.

For each security attribute, you can use the global security settings or customize settings for the domain.

External authorization providers:

Select to specify whether to use the default authorization configuration or an external authorization provider.

The external providers must be based on the Java uthorization Contract for Containers (JACC) specification to handle the Java(TM) 2 Platform, Enterprise Edition (J2EE) authorization. Do not modify any settings on the authorization provider panels unless you have configured an external security provider as a JACC authorization provider.

Custom properties:

Select to specify name-value pairs of data, where the name is a property key and the value is a string.

Specify extent of protection wizard settings

Use this security wizard page to determine whether to enable application security and restrict access to local resources. When you use the wizard, admin security is enabled by default.

To view this security wizard page, click **Security > Global security > Security configuration wizard**.

Enable application security:

Enables security for the applications in your environment. This type of security provides application isolation and requirements for authenticating application users

In previous releases of WebSphere Application Server, when a user enabled global security, both administrative and application security were enabled. In WebSphere Application Server Version 6.1, the previous notion of global security is split into administrative security and application security, each of which you can enable separately.

As a result of this split, WebSphere Application Server clients must know whether application security is disabled at the target server. Administrative security is enabled, by default. Application security is disabled, by default. To enable application security, you must enable administrative security. Application security is in effect only when administrative security is enabled.

Default: Disabled

Use Java 2 security to restrict application access to local resources:

Specifies whether to enable or disable Java 2 security permission checking. By default, access to local resources is not restricted. You can choose to disable Java 2 security, even when application security is enabled.

When the **Use Java 2 security to restrict application access to local resources** option is enabled and if an application requires more Java 2 security permissions than are granted in the default policy, the application might fail to run properly until the required permissions are granted in either the app.policy file or the was.policy file of the application. AccessControl exceptions are generated by applications that do not have all the required permissions. See the related links for more information about Java 2 security.

Default: Disabled

Security custom properties

Use this page to understand the predefined custom properties that are related to security.

To view this administrative console page, click **Security > Global security > Custom properties**. You can click **New** to add a new custom property and its associated value.

com.ibm.CSI.rmilnboundLoginConfig:

This property specifies the Java Authentication and Authorization Service (JAAS) login configuration that is used for Remote Method Invocation (RMI) requests that are received inbound.

By knowing the login configuration, you can plug in a custom login module that can handle specific cases for RMI logins.

Default `system.RMI_INBOUND`

com.ibm.CSI.rmiOutboundLoginConfig:

This property specifies the JAAS login configuration that is used for RMI requests that are sent outbound.

Primarily, this property prepares the propagated attributes in the Subject to be sent to the target server. However, you can plug in a custom login module to perform outbound mapping.

Default `system.RMI_OUTBOUND`

com.ibm.CSI.supportedTargetRealms:

This property enables credentials that are authenticated in the current realm to be sent to any realm that is specified in the Trusted target realms field. The Trusted target realms field is available on the CSiv2 outbound authentication panel. This property enables those realms to perform inbound mapping of the data from the current realm.

It is not recommended that you send authentication information to an unknown realm. Thus, this provides a way to specify that the alternate realms are trusted. To access the CSiv2 outbound authentication panel, complete the following steps:

1. Click **Security > Global security**.
2. Under RMI/IIOP security, click **CSiv2 outbound authentication**.

com.ibm.CSI.disablePropagationCallerList:

This property completely disables the caller list and will not allow the caller list to change. This property prevents the creation of multiple sessions.

This property completely disables adding a caller or host list in the propagation token. Setting this property can be a benefit when the caller or host list in the propagation token is not needed in the environment.

Note: If this property is set to **true** as well as **com.ibm.CSI.propagateFirstCallerOnly**, then **com.ibm.CSI.disablePropagationCallerList** takes precedence.

Default `false`

com.ibm.CSI.propagateFirstCallerOnly:

This property will not allow the caller list to change and thus prevent the creation of multiple session entries. This property specifically limits the caller list to the first caller only.

This property logs the first caller in the propagation token that stays on the thread when security attribute propagation is enabled. Without setting this property, all caller switches get logged, which affects performance. Typically, only the first caller is of interest.

Note: If this property is set to **true** as well as **com.ibm.CSI.disablePropagationCallerList**, then **com.ibm.CSI.disablePropagationCallerList** takes precedence.

Default false

com.ibm.security.useFIPS:

Specifies that Federal Information Processing Standard (FIPS) algorithms are used. The application server uses the IBMJCEFIPS cryptographic provider instead of the IBMJCE cryptographic provider.

Default false

com.ibm.websphere.security.krb.canonical_host:

This property specifies whether (true) or (false) the WebSphere Application Server uses the canonical form of the URL/HTTP host name in authenticating a client.

If this property is set to “false”, a Kerberos ticket can contain a host name that differs from the HTTP host name header. An error can occur as follows:

CWSPN0011E: An invalid SPNEGO token has been encountered while authenticating a HttpServletRequest

You can avoid an error message by setting this property to “true” and allowing WebSphere Application Server to authenticate using the canonical form of the URL/HTTP host name.

Default false

com.ibm.ws.security.createTokenSubjectForAsynchLogin:

In this release, the actual LTPA token data is not available from a WSCredential.getCredentialToken() call when called from an asynchronous bean. For an existing configuration, you can add the custom property com.ibm.ws.security.createTokenSubjectForAsynchLogin to allow the LTPAToken to be forwarded to asynchronous beans. This property allows portlets to successfully perform LTPA token forwarding.

Using the administrative console, create this custom property as follows:

1. Click **Security > Global security**.
2. Under Additional property, click **Custom properties**.
3. Click **New**.
4. In the Name field, type com.ibm.ws.security.createTokenSubjectForAsynchLogin

Note: This custom property name is case sensitive.

5. In the Value field, type true
6. Click **Apply** and **Save**, then restart the WebSphere Application Server.

Note: This custom property applies only to system conditions where Server A makes EJB calls from asynchronous beans to Server B. This property does not apply for JAAS login situations.

Default not applicable

com.ibm.ws.security.defaultLoginConfig:

This property is the JAAS login configuration that is used for logins that do not fall under the WEB_INBOUND, RMI_OUTBOUND, or RMI_INBOUND login configuration categories.

Internal authentication and protocols that do not have specific JAAS plug points call the system login configuration that is referenced by com.ibm.ws.security.defaultLoginConfig configuration.

Default system.DEFAULT

com.ibm.ws.security.ssoInteropModeEnabled:

This property determines whether to send LtpaToken2 and LtpaToken cookies in the response to a Web request (interoperable).

When this property value is false, the application server just sends the new LtpaToken2 cookie which is stronger, but not interoperable with some other products and Application Server releases prior to Version 5.1.1. In most cases, the old LtpaToken cookie is not needed and you can set this property to false.

Default true

com.ibm.ws.security.webChallengeCustomSubjectNotFound:

This property determines the behavior of a single sign-on LtpaToken2 login.

If the token contains a custom cache key and the custom Subject cannot be found, then the token is used to log in directly as the custom information needs to be regathered if this property value is set to true. A challenge also occurs so that the user is required to login again. When this property value is set to false and the custom Subject is not found, the LtpaToken2 is used to login and gather all of the registry attributes. However, the token might not obtain any of the special attributes that downstream applications might expect.

Default true

com.ibm.ws.security.webInboundLoginConfig:

This property is the JAAS login configuration that is used for Web requests that are received inbound.

By knowing the login configuration, you can plug in a custom login module that can handle specific cases for Web logins.

Default system.WEB_INBOUND

com.ibm.ws.security.webInboundPropagationEnabled:

This property determines whether a received LtpaToken2 cookie should search for the propagated attributes locally before searching the original login server that is specified in the token. After the propagated attributes are received, the Subject is regenerated and the custom attributes are preserved.

You can configure the data replication service (DRS) to send the propagated attributes to front-end servers such that a local dynacache lookup can find the propagated attributes. Otherwise, an MBean request is sent to the original login server to retrieve these attributes.

Default true

com.ibm.wsspi.security.ltpa.tokenFactory:

This property specifies the Lightweight Third Party Authentication (LTPA) token factories that can be used to validate the LTPA tokens.

Validation occurs in the order in which the token factories are specified because LTPA tokens do not have object identifiers (OIDs) that specify the token type. The Application Server validates the tokens using each token factory until validation is successful. The order that is specified for this property is the most likely order of the received tokens. Specify multiple token factories by separating them with a pipe (|) without spaces before or following the pipe.

Default `com.ibm.ws.security.ltpa.LTPATokenFactory |
com.ibm.ws.security.ltpa.LTPAToken2Factory |
com.ibm.ws.security.ltpa.AuthzPropTokenFactory`

com.ibm.wsspi.security.token.authenticationTokenFactory:

This property specifies the implementation that is used for an authentication token in the attribute propagation framework. The property provides an old LTPA token implementation for use as the authentication token.

Default `com.ibm.ws.security.ltpa.LTPATokenFactory`

com.ibm.wsspi.security.token.authorizationTokenFactory:

This property specifies the implementation that is used for an authorization token. This token factory encodes the authorization information.

Default `com.ibm.ws.security.ltpa.AuthzPropTokenFactory`

com.ibm.wsspi.security.token.propagationTokenFactory:

This property specifies the implementation that is used for a propagation token. This token factory encodes the propagation token information.

The propagation token is on the thread of execution and is not associated with any specific user Subjects. The token follows the invocation downstream wherever the process leads.

Default `com.ibm.ws.security.ltpa.AuthzPropTokenFactory`

com.ibm.wsspi.security.token.singleSignonTokenFactory:

This property specifies the implementation that is used for a Single Sign-on (SSO) token. This implementation is the cookie that is set when propagation is enabled regardless of the state of the `com.ibm.ws.security.ssoInteropModeEnabled` property.

By default, this implementation is the `LtpaToken2` cookie.

Default `com.ibm.ws.security.ltpa.LTPAToken2Factory`

security.enablePluggableAuthentication:

This property is no longer used. Instead, use `WEB_INBOUND` login configuration.

Complete the following steps to modify the `WEB_INBOUND` login configuration:

1. Click **Security > Global security**.
2. Under Java Authentication and Authorization Service, click **System logins**.

Default true

com.ibm.CSI.rmilnboundMappingConfig:

This property defines the system JAAS login configuration that is used to perform application specific principal mapping.

Default None

com.ibm.CSI.rmilnboundMappingEnabled:

This property, when set to true, enables the application specific principal mapping capability.

Default false

com.ibm.CSI.rmiOutboundMappingEnabled:

This property, when set to true, enables the original caller subject embedded in the WSSubjectWrapper object to be restored.

Default false

security.useDefaultPolicyWhenJ2SDisabled:

The NullDynamicPolicy.getPermissions method provides an option to delegate a default policy class to construct a Permissions object when this custom security is set to true. When this property is set to false, an empty Permissions object is returned.

Complete the following steps to set this property:

1. Log in to the administrative console.
2. Click **Security > Secure administration, applications, and infrastructure > Custom properties**.
3. Click **New** and add the following values:

Name field

security.useDefaultPolicyWhenJ2SDisabled

Value field

true

4. Click **OK**, then **Save**

Default false

com.ibm.websphere.lookupRegistryOnProcess:

This property can be set when realm registry lookups are performed via an MBean on a remote server if the realm is local OS security.

By default, the user registry tasks listRegistryUsers and listRegistryGroups perform lookups from the current process. In the case of Network Deployment (ND), that is the dmgr.

When dealing with a local OS user registry, lookup should occur on the actual server where the registry resides. In an ND environment that could be a remote machine. To perform lookup on the server process where the registry resides, the `com.ibm.websphere.lookupRegistryOnProcess` custom property should be set to `true`.

If `com.ibm.websphere.lookupRegistryOnProcess` is not set, or set to `false`, then the lookup is performed on the current process. The custom property can be set using the `setAdminActiveSecuritySettings` task for global security or the `setAppActiveSecuritySettings` task for a security domain.

Security custom property collection

Use this page to view and manage arbitrary name-value pairs of data, where the name is a property key and the value is a string value that can be used to set internal system configuration properties.

The administrative console contains several custom properties pages that work similarly. To view one of these administrative pages, click a **Custom properties** link.

Name:

Specifies the name (or key) for the property.

Each property name must be unique. If the same name is used for multiple properties, the value specified for the first property is used.

Do not start your property names with `was.` because this prefix is reserved for properties that are predefined in the application server.

Value:

Specifies the value paired with the specified name.

Description:

Provides information about the name-value pair.

Security custom property settings

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

The administrative console contains several custom property settings pages that work similarly. To view one of these administrative pages, click **Custom properties**.

Name:

Specifies the name (or key) for the property.

Each property name must be unique. If the same name is used for multiple properties, the value specified for the first property is used.

Do not start your property names with `was.` because this prefix is reserved for properties that are predefined in the product.

Data type String

Value:

Specifies the value paired with the specified name.

Data type String

Description:

Provides information about the name and value pair.

Data type String

Testing security after enabling it

Basic tests are available that show whether the fundamental security components are working properly. Use this task to validate your security configuration.

Before you begin

After configuring administrative security and restarting all of your servers in a secure mode, validate that security is properly enabled.

There are a few techniques that you can use to test the various security login types. For example, you can test the Web-based BasicAuth login, Web-based form login, and the Java client BasicAuth login.

Basic tests are available that show whether the fundamental security components are working properly. Complete the following steps to validate your security configuration:

1. After enabling security, verify that your system comes up in secure mode.
2. Test the Web-based BasicAuth with *Snoop*, by accessing the following URL: `http://hostname.domain:9080/snoop`.

A login panel is displayed. If a login panel does not display, then a problem exists. If the panel appears, type in any valid user ID and password in your configured user registry.

Note: The Snoop servlet is only available in the domain if you included the **DefaultApplication** option when adding the application server to the cell. The **-includeapps** option for the **addNode** command migrates the **DefaultApplication** option to the cell. Otherwise, skip this step.

3. Test the Web-based form login by starting the administrative console: `http://hostname.domain:port_number/ibm/console`. A form-based login page is displayed. If a login page does not appear, try accessing the administrative console by typing `https://myhost.domain:9043/ibm/console`.

Type in the administrative user ID and password that are used for configuring your user registry when configuring security.

4. Test Java Client BasicAuth with *dumpNameSpace*.

Use the `app_server_root/bin/dumpNameSpace.bat` file. A login panel appears. If a login panel does not appear, there is a problem. Type in any valid user ID and password in your configured user registry.

5. Test all of your applications in secure mode.
6. If all the tests pass, proceed with more rigorous testing of your secured applications. If you have any problems, review the SYSOUT and SYSPRINT logs. For more information on common problems, see Chapter 14, "Troubleshooting security configurations," on page 1289.

Results

The results of these tests, if successful, indicate that security is fully enabled and working properly.

The Security Configuration Wizard

The Security Configuration Wizard guides you through the process of completing the basic requirements to secure your application serving environment.

This wizard is available from the Security menu from the left pane of the admin console. To get to the wizard, navigate to **Security** → **Global security** → **Security Configuration Wizard**.

Step one of the configuration wizard allows you to choose the level of security desired. application-level security is selected by default. You also have the option of selecting Java 2 security.

Step two of the configuration wizard allows you to select a user repository. You have the following options:

- “Federated repository wizard settings” on page 177
- “Local operating system wizard settings” on page 119
- “Standalone custom registry wizard settings” on page 145
- “Standalone LDAP registry wizard settings” on page 125

Step three of the configuration wizard allows you to specify the local operating system user and group definitions as the repository, and, if necessary, to provide the name of a user with administrator privileges.

Step four of the configuration wizard provides a summary of the results of the configuration process.

Securing specific application servers

You can customize security to some extent at the application server level. You can disable administrative security on an application server.

Before you begin

Note: Server level security has been deprecated in this release of WebSphere Application Server. Multiple security domain support has been added in its place. You can create different security configurations and assign them to different applications in WebSphere Application Server processes. By creating multiple security domains, you can configure different security attributes for both administrative and user applications within a cell environment. You can configure different applications to use different security configurations by assigning the servers or clusters or SIBuses that host these applications to the security domains. Read about “Multiple security domains” on page 78 for more detailed information.

You can also modify Java 2 Security and some of the other security attributes that are found on the Global security panel. This panel provides access to the cell-level security settings. You cannot configure a different authentication mechanism or user registry on an individual server basis. This feature is limited to cell-level configuration only.

By default, server security inherits all of the values that are configured for cell-level security. To override the cell-level security configuration at the server level, click **Servers > Application Servers > server_name**. Under Security, click **Server Security** and click any of the following links:

- CSiv2 inbound authentication
- CSiv2 outbound authentication
- CSiv2 inbound transport
- CSiv2 outbound transport
- SAS inbound transport
- SAS outbound transport
- Server-level security

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

After modifying the configuration in any of these panels and clicking **OK** or **Apply**, the security configuration for that panel or set of panels now overrides cell-level security. Other panels that are not overridden continue to be inherited at the cell-level. However, you can always revert back to the cell-level configuration at any time. You can revert back to the cell-level security configuration by clearing the check box next to any of the following options on the Server security panel:

- Security settings for this server override cell settings
- RMI/IIOP security for this server overrides cell settings
- SAS security for this server overrides cell settings

For more information, see “Server and administrative security” on page 41.

1. Start the administrative console for the deployment manager. To get to the administrative console, go to `http://host.domain:port_number/ibm/console`. If security is disabled, you can enter any ID. If security is enabled, you must enter a valid user ID and password, which is either the administrative ID that is configured for the user registry or a user ID that is entered as an administrative user. To add a user ID as an administrative user, click **System Administration > Console settings > Console users**.
2. Configure cell-level security if you have not configured it previously. Go to “Enabling security” on page 37 for detailed steps. After security is configured, configure server-level security.

Note: Server-level security is not enabled when you select the Enable application security option on the Server-level security settings of the administrative console. You also must enable cell-level security by selecting the Enable administrative security option on the Global security settings panel of the administrative console.

3. To configure server-level security, click **Servers > Application Servers > server name**. Under Security, click **Server security**. The status of the security level that is in use for this application server is displayed.

By default, you can see that your cell-level security configuration, Common Secure Interoperability (CSI), and SAS have not been overridden at the server level. CSI and SAS are authentication protocols for RMI/IIOP security requests. The server-level security panel lists attributes that are on the Global security panel and can be overridden at the server level. Not all of the attributes on the Global security panel can be overridden at the server level, including the user account repository.

4. To enable administrative security for this application server, go to the Server-level security panel, select the **Security settings for this server override cell settings** and the **Enable application security** options. By modifying the Server-level security panel, these settings override the settings for cell-level security.
5. Click **Apply** and **Save**.
6. To enable RMI/IIOP security for the application server, go to the Server-level security panel, select the **RMI/IIOP security for this server overrides cell settings** option and click **Apply**. If you select the **RMI/IIOP security for this server overrides cell settings** option, any changes that you make to the CSIv2 authentication or transport settings override the same settings on the cell level.

What to do next

Typically, server-level security is used to disable user security for a specific application server. However, this can also be used to disable or enable the Java 2 security manager, and to configure the authentication requirements for RMI/IIOP requests both incoming and outgoing from this application server.

After you modify the configuration for a particular application server, you must restart the application server for the changes to become effective. To restart the application server, go to **Servers > Application servers** and click the server name that you recently modified. Click **Stop** and then **Start**.

If you disabled security for the application server, you can typically test a Web address that is protected when security is enabled.

One URL that usually is installed when the DefaultApplication during installation is the snoop application. If the DefaultApplication is installed on the application server, test that security is disabled by going to the following URL: `http://host.domain:9080/snoop`. If security is disabled, a prompt does not display. This URL is just one method of validating the configuration. Validate that the configuration is appropriate for your applications.

Server-level security settings

Use this page to enable server-level security and specify other server-level security configurations.

Note: Server level security has been deprecated in this release of WebSphere Application Server. Multiple security domain support has been added in its place. You can create different security configurations and assign them to different applications in WebSphere® Application Server processes. By creating multiple security domains, you can configure different security attributes for both administrative and user applications within a cell environment. You can configure different applications to use different security configurations by assigning the servers or clusters or SIBuses that host these applications to the security domains. Read about “Multiple security domains” on page 78 for more detailed information.

To view this administrative console page, complete the following steps:

1. Click **Servers > Server Types > WebSphere application servers > *server_name***.
2. Click **Security > Security domain**.

Security settings for this server override cell settings

Specifies that the settings on this panel override the settings on the Security > Security applications, administration, and infrastructure panel.

Enable application security

Specifies that server security is disabled, by default, when application security is enabled. Administrative (administrative console and wsadmin) and naming security remain enabled while application security is enabled, regardless of the status of this flag.

Default Disabled

Use Java 2 security to restrict application access to local resources

Specifies that the server enforces Java 2 security permission checking at the server level. When cleared, the Java 2 server-level security manager is not installed and all of the Java 2 security permission checking is disabled at the server level.

If your application policy file is not set up correctly, see “Configuring the was.policy file” on page 750 in the information center.

Default Disabled

Warn if applications are granted custom permissions

Specifies whether a warning is issued during application installation when an application requires a Java 2 permission that is normally not granted to an application.

The application server provides support for policy file management. A number of policy files are included in the application server. Some of these policy files are static and some of them are dynamic. Dynamic policy

is a template of permissions for a particular type of resource. In dynamic policy files, the code bases are evaluated at runtime using the configuration data. You can add or remove permissions, as needed, for each code base.

However, do not add, remove, or modify the existing code bases. The real code base is dynamically created from the configuration and run-time data. The `filter.policy` file contains a list of permissions that an application does not have, according to the J2EE 1.4 Specification. For more information on permissions, see the documentation on the Java 2 security policy files.

Default Enabled

Restrict access to resource authentication data

Enable this option to restrict application access to sensitive Java Connector Architecture (JCA) mapping authentication data.

Default Disabled

Use domain-qualified user names

Specifies whether user IDs that are returned by `getUserPrincipal()` are qualified with the server level security domain within which they reside.

Default Disabled

Authentication cache timeout

Specifies the time period in which a Lightweight Third Party Authentication (LTPA) token expires.

Data type	Integer
Units	Minutes and seconds
Default	10 minutes and 0 seconds
Range	Greater than 30 seconds. Avoid setting Authentication cache timeout value to 30 seconds or less.

Active protocol

Specifies the active server level security authentication protocol when server level security is enabled.

You can use an Object Management Group (OMG) protocol called Common Secure Interoperability Version 2 (CSIv2) for more vendor interoperability and additional features. If all of the servers in your entire security domain are Version 5.0 servers, it is best to specify CSI as your protocol.

This field displays if a version 6.0.x server exists in your environment.

Data type	String
Default	CSI and SAS
Range	CSI, CSI and SAS

RMI/IIOP security for this server overrides cell settings

Specifies that the Remote Method Invocation over Internet InterORB Protocol (RMI/IIOP) settings on this panel override the settings on the Security > Security applications, administration, and infrastructure panel.

Default Enabled

Chapter 4. Configuring multiple security domains

By default, all administrative and user applications in WebSphere Application Server use the global security configuration. For example, a user registry defined in global security is used to authenticate users for every application in the cell. Out-of-the-box, this behavior is the same as it was in previous releases of WebSphere Application Server. You can create additional WebSphere security domains if you want to specify different security attributes for some or all of your user applications. This section describes how to configure a security domain by using the administrative console.

Before you begin

Only users assigned to the administrator role can configure or create new multiple security domains. Enable global security in your environment before configuring multiple security domains.

Read about “Multiple security domains” on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

About this task

Security domains enable you to define multiple security configurations for use in your environment. For example, you can define different security (such as a different user registry) for user applications than for administrative applications. You can also define separate security configurations for user applications deployed to different servers and clusters.

Perform the following steps to configure a new security domain by using the administrative console:

1. Click **Security > Security domains**.
2. If you are creating a new multiple security domain, click **New**. Supply a unique name and description for the domain and click **Apply**. If you want to configure an existing multiple security domain, select one to edit. Once you click **Apply** the domain name and additional sections are displayed. One section enables you to define the security attributes for the domain, and another section enables you to select the scopes to which the domain applies.
3. Under Assigned Scopes, select whether you want to assign the security domain to the entire cell or if you want to select the specific servers, clusters, and service integration buses to be included in the security domain. The Assigned Scopes section has two views. The default view is a cell topology. To assign the security domain to the entire cell, click the check box for the cell and then click **Apply** or **OK**.

The name of the security domain appears next to the cell name, which indicates that the domain is now assigned to the cell. You can expand the topology and assign the domain to one or more servers and clusters. When an item in the topology is already assigned to another security domain, the check box is disabled and the name of the assigned domain is displayed to the right of the scope name. If you want to assign one of these scopes to the domain, you must first disassociate it with its current domain.

Select **All assigned scopes** to view a list of only those resources that are currently assigned to the security domain.

4. Customize your security configuration by specifying security attributes for your new domain. Attributes that are not listed can not be customized at the domain level. Domains inherit attributes from the global security configuration.

There are twelve individually configurable security attribute sections. You can expand and collapse each section. In the collapsed state, the name and a summary value for the section are displayed. Additionally, the summary value text indicates whether the attribute is defined in global security and is reused by the domain (as indicated by gray text) or if it is customized for the domain (as indicated by black text prefixed by the word “Customized”).

Initially, each security attribute is set to use the global security settings. When an attribute is set to use global security, there is no domain-specific configuration for that attribute. Applications that use the domain use the global configuration for these security attributes.

Only configure the security attributes that you want to change. To configure a security attribute for a domain, expand the security attribute section. The key properties of the global configuration display beneath the **Use global security** option. These properties are provided for convenience.

To customize the configuration for the domain, select **Customize for this domain**. Configure the property and then click **OK** or **Apply**.

Note: In general, when you select **Customize for this domain**, you override all of the security configurations that are defined for that section in global security. Application logins, system logins, and J2C authentication data entries are some exceptions. When you define entries for a domain, applications in the domain are able to access the global entries in addition to the domain-specific entries.

For example, you might want to use a different user registry for applications that use the security domain but also want to use the global security configuration for all of the other security properties. In this case, expand the User Realm section and select **Customize for this domain**. Select a user registry type, click **Configure**, and provide the appropriate configuration details on the subsequent panel.

You can change security attributes such as the following:

Application Security

Specifies the settings for application security and Java 2 security. You can use the global security settings or customize the settings for a domain.

Select **Enable application security** to enable or disable security this choice for user applications. When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

Java 2 Security

Select **Java 2 security** to enable or disable Java 2 security at the domain level. This choice enables or disables Java 2 security at the process (JVM) level so that all applications (both administrative and user) can enable or disable Java 2 security.

User realm

This section enables you to configure the user registry for the security domain. You can separately configure any registry except the federated registry that is used at the domain level. The federated repository can only be configured at the global level but can be used at the domain level. Read about “Multiple security domains” on page 78 for more information.

Trust association

When you configure the trust association interceptor (TAI) at a domain level, the interceptors configured at the global level are copied to the domain level for convenience. You can modify the interceptor list at the domain level to fit your needs. Only configure those interceptors that are to be used at the domain level.

SPNEGO Web Authentication

The SPNEGO Web authentication, which enables you to configure SPNEGO for Web resource authentication, can be configured at the domain level.

Note: In WebSphere Application Server Version 6.1, a TAI that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server

7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

RMI/IOP Security

The RMI/IOP security attribute refers to the CSv2 (Common Secure Interoperability version 2) protocol properties. When you configure these attributes at the domain level, the RMI/IOP security configuration at the global level is copied for convenience.

You can change the attributes that need to be different at the domain level. The Transport layer settings for CSv2 inbound communications should be the same for both the global and the domain levels. If they are different, the domain level attributes are applied to all of the application in the process.

JAAS application logins

Specifies the configuration settings for the Java Authentication and Authorization Service (JAAS) application logins. You can use the global security settings or customize the settings for a domain.

Note: The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

JAAS system logins

Specifies the configuration settings for the JAAS system logins. You can use the global security settings or customize the configuration settings for a domain.

JAAS J2C authentication

Specifies the configuration settings for the JAAS J2C authentication data. You can use the global security settings or customize the settings for a domain.

Authentication Mechanism Attributes

Specifies the various cache settings that need to be applied at the domain level.

Select **Authentication cache settings** to specify your authentication cache settings. The configuration specified on this panel is applied only to this domain.

Select **LTPA Timeout** to configure a different LTPA timeout value at the domain level. The default timeout value is 120 minutes, which is set at the global level. If the LTPA timeout is set at the domain level, any token that is created in the security domain when accessing user applications is created with this expiration time.

When **Use realm-qualified user names** is enabled, user names returned by methods such as `getUserPrincipal()` are qualified with the security realm (user registry) used by applications in the security domain.

Authorization Provider

You can configure an external third party JACC (Java Authorization Contract for Containers) provider at the domain level. Tivoli Access Manager's JACC provider can only be configured at the global level. Security domains can still use it if they do not override the authorization provider with another JACC provider or with the built-in native authorization.

Custom properties

Set custom properties at the domain level that are either new or different from those at the global level. By default, all of the custom properties at the global security configuration can be

accessed by all of the applications in the cell. The security runtime code first checks for the custom property at the domain level. If it does not find it, it then attempts to obtain the custom property from the global security configuration.

5. Once you have configured the security attributes and assigned the domain to one or more scopes, click **Apply** or **OK**.
6. Restart all servers and clusters for your changes to take effect.

Multiple security domains

The WebSphere Security Domains (WSD) provide the flexibility to use different security configurations in WebSphere Application Server. The WSD is also referred to as multiple security domains, or simply, security domains. You can configure different security attributes, such as the UserRegistry, for different applications.

Note: Multiple security domain support is new in this release of WebSphere Application Server. You can create different security configurations and assign them to different applications in WebSphere Application Server processes. By creating multiple security domains, you can configure different security attributes for both administrative and user applications within a cell environment. You can configure different applications to use different security configurations by assigning the servers or clusters or service integration buses that host these applications to the security domains. Only users assigned to the administrator role can configure multiple security domains.

The following sections describe multiple security domains in more detail:

- “Why security domains are useful”
- “Relationship between global security and security domains” on page 79
- “Contents of a security domain” on page 80
- “Creating security domains” on page 81
- “Configuring attributes for security domains” on page 82
- “Associating scopes to security domains” on page 83
- “Relationship between old server level security and the new security domains” on page 84
- “How domain level security attributes are used by security runtime and applications” on page 84
- “Client and application security programming model when using security domains” on page 86
- “Application deployment in multiple domains configurations” on page 88
- “Cross realm communication” on page 89
- “Federating a node with security domains” on page 91
- “Security domains in a mixed-version environment” on page 92
- “Modifying security domains” on page 92

Why security domains are useful

WebSphere Security Domains provide two major benefits:

- WebSphere Application Server administrative applications can be configured with a different set of security configurations from those for user applications.
- They enable one set of applications to have a different set of security configurations from another set of applications.

For example, WebSphere Application Server administration can be configured to a user registry of federated repository while the applications can be configured to a user registry of LDAP.

In previous versions of WebSphere Application Server, all administrative and user applications use security attributes different from those attributes that are defined in global security. All administrative and user

applications in WebSphere Application Server use global security attributes by default. For example, a user registry defined in global security is used to authenticate a user for every application in the cell.

In this release of WebSphere Application Server, however, you can use multiple security attributes for user applications other than the global security attributes, create a security domain for those security attributes that must differ, and associate them with the servers and clusters that host those user applications. You also can associate a security domain with the cell. All of the user applications in the cell use this security domain if they do not have a domain previously associated with them. However, global security attributes are still required for administrative applications such as the administrative console, naming resources and MBeans.

If you have used server level security in previous releases of WebSphere Application Server, you should now use multiple security domains since they are more flexible and easier to configure.

Server level security is deprecated in this release. Read “Relationship between global security and security domains” for more information.

Relationship between global security and security domains

Global Security applies to all administrative functions and the default security configuration for user applications. Security domains can be used to define a customized configuration for user applications.

You must have a global security configuration defined before you can create security domains. The global security configuration is used by all of the administrative applications such as the administrative console, naming resources, and Mbeans. If no security domains are configured, all of the applications use information from the global security configuration. User applications such as Enterprise JavaBeans (EJBs), servlets and administrative applications use the same security configuration.

When you create a security domain and associate it with a scope, only the user applications in that scope use the security attributes that are defined in the security domain. The administrative applications as well as the naming operations in that scope use the global security configuration. Define the security attributes at the domain level that need to be different from those at the global level. If the information is common, the security domain does not need to have the information duplicated in it. Any attributes that are missing in the domain are obtained from the global configuration. The global security configuration data is stored in the `security.xml` file, which is located in the `$WAS_HOME/profiles/$ProfileName/cells/$CellName` directory.

The following figure provides an example of a security multiple domain where the cell, a server and a cluster are associated with different security domains. As shown in the figure, the user applications in server S1.1 as well as the cluster use security attributes that are defined in Domain2 and Domain3 respectively (since these scopes are associated with these domains). Server S2.2 is not associated with a domain. As a result, the user application in S2.2 uses the domain that is associated with the cell (Domain1) by default . Security attributes that are missing from the domain level are obtained from the global configuration.

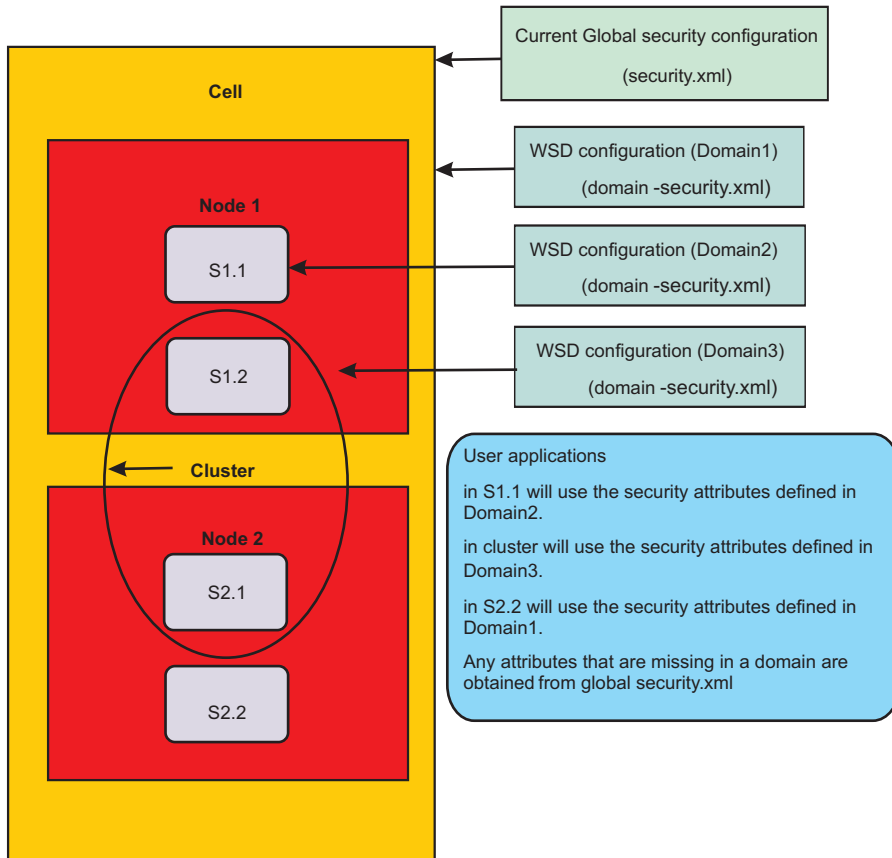


Figure 1. Scopes that can be associated to a security domain

The following figure shows the various high-level security attributes that can be defined at the global security configuration and those that can be overridden at the domain level.




Figure 2. Security attributes that can be configured at the security domain

Contents of a security domain

A security domain is represented by two configuration files. One configuration file contains the list of attributes that are configured in the security domain. The other configuration file contains the scopes that

use the security domain. The security domain information is stored in the `$WAS_HOME/profiles/$ProfileName/config/waspolicies/default/securitydomains/$SecurityDomainName` directory. For every security domain that is configured, a `$SecurityDomainName` directory is created with two files in it: the `security-domain.xml` file contains the list of security attributes configured for the security domain, and the `security-domain-map.xml` file contains the scopes that use the security domain.

The following figure indicates the location of the main security domain related files and the contents of those files.

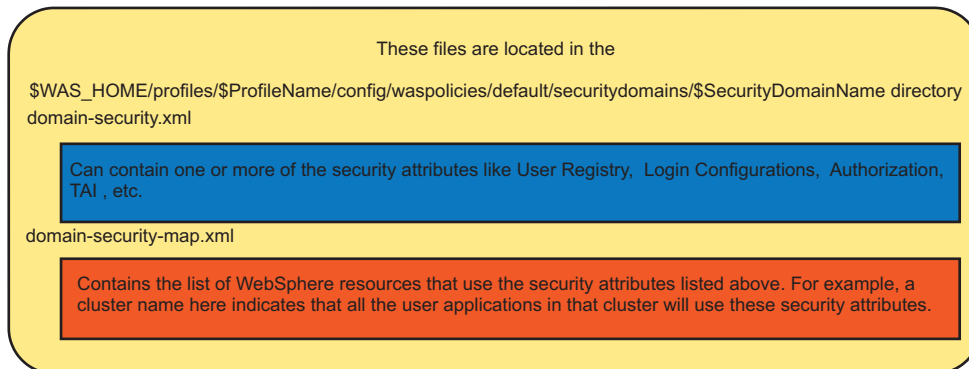


Figure 3. Location and contents of the main security domain related files

Note: You should not modify these files manually. Use administrative console tasks or scripting commands to modify the files instead. For a complete list of administrative tasks and scripting commands, see the links in "Related tasks" at the bottom of this document.

Creating security domains

Use the administrative console tasks or scripting commands to create security domains. In the administrative console, access security domains by clicking **Security > Security domains**. Help is available for each administrative console panel.

For a complete list of administrative console tasks and scripting commands, see the links in "Related tasks" at the bottom of this document.

When you create a security domain you must supply a unique name for the domain, the security attributes you want to configure for the security domain, and the scopes that need to use the security domain. Once configured, the servers that use the security domain must be restarted. The user applications in those scopes then use the attributes that are defined in the security domain. Any attributes that are not configured at the domain level are obtained from the global security configuration. Administrative applications and naming operations in the scopes always use the security attributes from the global security configuration. You must actively manage these attributes.

Any new security domain attributes must be compatible with those global security attributes that are inherited by the user applications that are assigned to the domain.

Other than for JAAS and custom properties, once global attributes are customized for a domain they are no longer used by user applications.

The security domains panel in the administrative console enables you to assign resources and to select the appropriate security attributes for your domain. The panel displays the key security attributes at the global configuration; you can make the decision to override them at the domain level if necessary. Once

you have configured and saved the attributes at the domain level, the summary value on the panel displays the customized value for the domain (tagged with the word "customized" in black text).

A scope (a server, cluster, SIBus or a cell) can be associated with only one domain. For example, you cannot define two domains that both have the cell-wide scope. Multiple scopes, however, can be defined in the same security domain. For example, a domain can be scoped to Server1 and to Server2 only within the cell.

The assigned scopes section on the security domain panel displays two views: one view that enables you to select and assign scopes to the domain, and another view that enables you to see a list of the currently assigned scopes. For convenience, you also have the flexibility to copy all of the security attributes from an existing security domain or the global configuration into a new security domain, and then modify only those attributes that must be different. You must still associate the scopes to these copied domains.

Scripting commands also provide you with the ability to create, copy and modify security domains. Once you create a domain, you must run the appropriate commands to associate security attributes and scopes to it.

Configuring attributes for security domains

Security attributes that can be configured at the domain level in WebSphere Application Server Version 7.0 are:

- Application security
- Java 2 security
- User realm (registry)
- Trust association
- Simple and Protected GSS-API Negotiation (SPNEGO) Web authentication
- RMI/IIOP security (CSIV2)
- JAAS logins (Application, System and J2C Authentication Data)
- Authentication mechanism attributes
- Authorization provider
- Custom properties

The security domains panels in the administrative console display all of these security attributes.

Some of the other well-known attributes that you cannot override at the domain level are Kerberos, Audit, Web Single Sign-on (SSO), Tivoli Access Manager (TAM), and federated repositories. The Secure Socket Layer (SSL) attribute already supports different scopes, but it is not part of the domain configuration. For all of the attributes that are not supported at the domain level, user applications in a domain share their configuration from the global level.

Any new security domain attributes must be compatible with those global security attributes that are inherited by the user applications that are assigned to the domain. You must actively manage these attributes. For example, if you customize only a JAAS configuration at the domain level you must make sure that it works with the user registry configured at the global level (if the user registry is not customized at the domain level).

Other than for JAAS and custom properties, once global attributes are customized for a domain they are no longer used by user applications.

Tivoli Access Manager is used to provide authentication (TrustAssociationInterceptor, PDLoginModule) and authorization (JACC). You can configure these Tivoli Access Manager attributes at the global level. You cannot have a different Tivoli Access Manager configuration at the domain level to override the

configuration at the global level. However, you can override these Tivoli Access Manager attributes at the domain level with different Authentication and Authorization security attributes. For example, you can use a different Trust Association Interceptor (TAI) or a different authorization provider.

The federated repositories user registry can only be configured at the global level. There can only be one instance of federated repositories in WebSphere Application Server. Any security domain can use this configuration for its user registry. A security domain can override the configuration with its own user registry. However, it cannot configure a different instance of the federated repository.

Associating scopes to security domains

In WebSphere Application Server Version 7.0, you can associate a security domain at the cell level, the server level, the cluster level and the SIBus level.

Note: For more information about SIBus and bus security in multiple security domains for WebSphere Application Server Version 7.0, read about "Bus security in multiple security domains" in the InfoCenter.

When a security domain is associated with a server that is not part of a cluster, all user applications in that server use the attributes from the security domain. Any missing security attributes are obtained from the global security configuration. If the server is part of a cluster, you can associate the security domain with the cluster but not with the individual members in that cluster. The security behavior then remains consistent across all of the cluster members.

If a server is to be part of a cluster, create a cluster first and associate the security domain to it. You might have associated a domain to a server before it was a member of a cluster. If so, even though the domain is associated with the server directly, the security runtime code does not look at the domain. When a server is a cluster member, the security runtime disregards any security domains associated directly to the server. Remove the server scope from the security domain and associate the cluster scope to it instead.

A security domain can also be associated to the cell. This is usually done when you want to associate all user applications in WebSphere Application Server to a security domain. In this scenario, all of the administrative applications and the naming operations use the global security configuration while all of the user applications use the domain level configuration. If you want to split the security configuration information for administrative and user applications, this is all that is needed.

If you have a mixed-version environment, or plan to have one in future, and you want to associate security domains at the cell level, read "Security domains in a mixed-version environment" on page 92 for more information.

If you are on a base profile server that has its own security domain defined, which is then federated to a deployment manager, associate the server scope to the security domain and not the cell scope. When you federate that node, the security domain information is propagated to the deployment manager. If the cell scope is associated to it, the network deployment configuration uses this security configuration, which might impact existing applications. During federation, the cell scope is changed to the server scope that is being federated. If the server scope is associated with the security domain, only that server uses the security domain after the federation. Other applications in other servers and clusters are not impacted. However, if this base profile server is registered to the Administrative Agent process you can associate the cell scope to the security domain if you want all of the servers from the base profile to use the same security domain for all of their user applications. Read about "Federating a node with security domains" on page 91 for more information.

You can have a security domain associated at the cell level and also other security domains associated to various clusters or individual servers (those that are not part of any clusters). In this case, the security runtime first checks if any security domains are associated with the server or a cluster. If there is a security domain associated with the server or a cluster, the security attributes defined in it are used for all

of the applications in that server or cluster. Any security attributes missing from this server or cluster domain are obtained from the global security configuration, and not from the domain configuration associated with the cell.

If the server or cluster does not have its own domain defined, the security runtime code uses the security attributes from the domain associated with the cell (if one is defined). Any security attributes missing from the cell domain are inherited from the global security configuration.

Relationship between old server level security and the new security domains

In previous releases of WebSphere Application Server, you could associate a small set of security attributes at a server level. These attributes were used by all of the applications at the server level. The previous way of configuring the security attributes is deprecated in WebSphere Application Server 7.0, and will be removed in a future release.

You should now use the new security domains support in WebSphere Application Server 7.0, as these security domains are more easily managed and much more flexible. For example, in previous versions of WebSphere Application Server, you must manually associate the same security configuration to all of the cluster members by configuring the same security attributes for every server in a cluster.

The migration tool migrates the existing server level security configuration information to the new security domain configuration when the script compatibility mode is `false` (`-scriptCompatibility="false"`). A new security domain is created for every server security configuration if it is not part of a cluster. If it is part of a cluster, a security domain is associated with the cluster instead of with all of the servers in that cluster. In both cases, all of the security attributes that were configured at the server level in previous releases are migrated to the new security domain configuration, and the appropriate scope is assigned to the security domains.

If the script compatibility mode is set to `true`, the server level security configuration is not migrated to the new security domains configuration. The old server security configuration is migrated without any changes. The security runtime detects that the old security configuration exists and uses that information, even if a security domain is associated either directly or indirectly to the server. If the script compatibility mode is set to `true`, remove the security configuration from the server level and then create a security domain with the same set of security attributes.

How domain level security attributes are used by security runtime and applications

This section describes how the individual attributes at the domain level are used by the security runtime and how that impacts the user application security. Since all of these security attributes are also defined at the global level, more information about these attributes can be obtained elsewhere. For the purposes of this section, the emphasis is on domain level behavior.

1. Application Security:

Select **Enable application security** to enable or disable security for user applications. When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

2. Java 2 Security:

Select **Use Java 2 security** to enable or disable Java 2 security at the domain level or to assign or add properties related to Java 2 security. This choice enables or disables Java 2 security at the process (JVM) level so that all applications (both administrative and user) can enable or disable Java 2 security.

3. **User Realm (User Registry):**

This section enables you to configure the user registry for the security domain. You can separately configure any registry except the federated registry that is used at the domain level. The federated repository can only be configured at the global level but can be used at the domain level. Read about “Configuring attributes for security domains” on page 82 for more information.

When configuring a registry at the domain level you can choose to define your own realm name for the registry. The realm name distinguishes one user registry from another. The realm name is used in multiple places – in the Java client login panel to prompt the user, in the authentication cache, and when using native authorization.

At the global configuration level, the system creates the realm for the user registry. In previous releases of WebSphere Application Server, only one user registry is configured in the system. When you have multiple security domains you can configure multiple registries in the system. For the realms to be unique in these domains, configure your own realm name for a security domain. You also can choose the system to create a unique realm name if it is certain to be unique. In the latter case, the realm name is based on the registry that is being used.

For LDAP registries, the host:port of the LDAP server is the system-generated realm name. For localOS, the name of the localOS machine is the realm name. For custom user registries, the realm is the one returned by the `getRealm ()` method of the custom registry implementation.

If the system generated realm names are not unique enough, you can choose the option for the system to generate the realm name. If not, choose a unique realm name for each security domain where you have the user registry configured. If the underlying user repository is the same, use the same realm name in different domains. From a security runtime perspective, same realm names have the same set of users and groups information. For example, when users and groups information is required from a realm, the first user repository that matches the realm is used.

When more than one user registry is in a process, the naming lookup that uses “UserRegistry” as the lookup name returns the user registry that is used by user applications. The user registry used by administrative applications is bound by the lookup name, “AdminUserRegistry”.

As described in “Cross realm communication” on page 89, when an application in one realm communicates with an application in another realm using LTPA tokens, the realms have to be trusted. The trust relationship can be established using the **Trusted authentication realms – inbound** link in the user registry panel or by using the **addTrustedRealms** command. You can establish trust between different realms. A user logged into one realm can access resources in another realm. If no trust is established between the two realms the LTPA token validation fails.

4. **Trust Association:**

When you configure the trust association interceptor (TAI) at a domain level, the interceptors configured at the global level are copied to the domain level for convenience. You can modify the interceptor list at the domain level to fit your needs. Only configure those interceptors that are to be used at the domain level.

Tivoli Access Manager’s trust association interceptors can only be configured at the global level. The domain configuration can also use them, but cannot have a different version of the trust association interceptor. Only one instance of Tivoli Access Manager’s trust association interceptors can exist in the cell.

5. **SPNEGO Web Authentication:**

The SPNEGO Web authentication, which enables you to configure SPNEGO for Web resource authentication, can be configured at the domain level.

Note: In WebSphere Application Server Version 6.1, a TAI that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

6. **RMI/IIOP Security (CSlv2):**

The RMI/IIOP security attribute refers to the CSiv2 (Common Secure Interoperability version 2) protocol properties. When you configure these attributes at the domain level, the RMI/IIOP security configuration at the global level is copied for convenience.

You can change the attributes that need to be different at the domain level. The Transport layer settings for CSiv2 inbound communications should be the same for both the global and the domain levels. If they are different, the domain level attributes are applied to all of the application in the process.

When a process communicates with another process with a different realm, the LTPA authentication and the propagation tokens are not propagated to the downstream server unless that server is listed in the outbound trusted realms list. This can be done using the **Trusted authentication realms – outbound** link on the **CSiv2 outbound communication** panel, or by using the **addTrustedRealms** command task. Read about “Cross realm communication” on page 89 for more information.

7. **JAAS (Java Authentication and Authorization Service):**

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

For JAAS and custom properties only, once global attributes are customized for a domain they can still be used by user applications.

8. **Authentication Mechanism Attributes:**

Specifies the various cache settings that must be applied at the domain level.

- a. Authentication cache settings - use to specify your authentication cache settings. The configuration specified on this panel is applied only to this domain.
- b. LTPA Timeout - You can configure a different LTPA timeout value at the domain level. The default timeout value is 120 minutes, which is set at the global level. If the LTPA timeout is set at the domain level, any token that is created in the security domain when accessing user applications is created with this expiration time.
- c. Use realm-qualified user names - When this selection is enabled, user names returned by methods such as `getUserPrincipal()` are qualified with the security realm (user registry) used by applications in the security domain.

9. **Authorization Provider:**

You can configure an external third party JACC (Java Authorization Contract for Containers) provider at the domain level. Tivoli Access Manager’s JACC provider can only be configured at the global level. Security domains can still use it if they do not override the authorization provider with another JACC provider or with the built-in native authorization.

10. **Custom properties:**

Set custom properties at the domain level that are either new or different from those at the global level. By default, all of the custom properties at the global security configuration can be accessed by all of the applications in the cell. The security runtime code first checks for the custom property at the domain level. If it does not find it, it then attempts to obtain the custom property from the global security configuration.

For JAAS and custom properties only, once global attributes are customized for a domain they can still be used by user applications.

Client and application security programming model when using security domains

A Java client or an application acting as a client that accesses an EJB typically does a naming lookup first. The naming resource, which is used by both administrative and the user applications, is considered an administrative resource. It is protected by the global security configuration information. In a multiple domain setup where the global security is using one realm (the *user registry*) and a domain is using a

different realm, the Java client must authenticate to two different realms. The first authentication is required for the realm in the global security configuration for the naming operation to succeed, and the second authentication is required to access the EJB, which uses a different realm.

The `CosNamingRead` role protects all naming read operations. This role is usually assigned the **Everyone** special subject. This implies that any user, valid or not, can look up the name space. When a multiple domain is defined, if the `CosNamingRead` role has the **Everyone** special subject the security runtime code in the client side does not prompt you to log in. It uses the `UNAUTHENTICATED` subject to access the naming operation instead. Once the naming lookup operation is completed, when the client attempts to access the EJB it is prompted with a login panel that indicates the realm that is currently used by that EJB application (that is, the realm used in the domain). The client then presents the appropriate user credentials for that realm, which can then access the EJB. This logic applies to all variations of login source, including `properties` and `stdin`, not just when the login source is set to `prompt`.

If the **Everyone** special subject is removed from the `CosNamingRead` role, you are prompted twice. If the login source is `properties`, you can uncomment the `com.ibm.CORBA.loginRealm` property in the `$WAS_HOME/profiles/$ProfileName/properties/sas.client.props` file and add the appropriate realms using “|” as the separator. You must also enter the corresponding users and passwords in the `com.ibm.CORBA.loginUserid` and `com.ibm.CORBA.loginPassword` properties respectively. When you are using the programmatic logon in the Java client code you must authenticate twice with different user credentials; once prior to do a naming lookup for the EJB (the user should be in the global realm), and later prior to calling any method in the EJB (the user should be in the EJB domain’s realm).

In general, when a Java client needs to authenticate to multiple and different realms it has to provide the credential information for all of those realms. If the login source is `prompt` or `stdin` it is prompted to login multiple times, once for each realm. If the login source is set to `properties`, the appropriate properties in the **sas.client.props** file (or any related file) are used for authenticating to different realms.

In certain scenarios, a client might make multiple calls to the same realm. For example, the Java client can access a resource using `realm1` followed by access to a resource using `realm2`, and then come back to access a resource in `realm1` again. In this case, the client is prompted three times; first for `realm1`, secondly for `realm2` and finally for `realm1` again.

By default, the subject that is used to login at a realm is not cached by the client side code. If you have this scenario, and you want the client to cache the subject based on the realm, set the `com.ibm.CSI.isRealmSubjectLookupEnabled` property to `true` in the **sas.client.props** file. If the `com.ibm.CSI.isRealmSubjectLookupEnabled` property is set, the client code caches the subject based on the realm name. The next time the Java client needs to authenticate to this realm, the cache is located to obtain the subject and the client is not prompted. Also, when the `com.ibm.CSI.isRealmSubjectLookupEnabled` property is set, the same subject that was logged in the first time is used for subsequent logins. If the subject information needs to change then this property should not be set.

If the client is doing a programmatic login it can pass the realm along with the user and password that it needs to authenticate. In this case, when the `com.ibm.CORBA.validateBasicAuth` property is set to `true` (the default value) in the **sas.client.props** file, the registry that matches the realm name is used for login. That realm must be supported in the process where the authentication takes place.

When using the `WSLogin` JAAS configurations, you also must set the `user_realm_callback` option in the **wsjaas_client.config** file in `$WAS_HOME/profiles/$ProfileName/properties` for the realm name to be passed to the call back handler. If you want to specify a different provider URL for the name server, set the `use_appcontext_callback` option and pass in the provider URL properties in a hash map to `WSLogin`.

If you do not know the realm name, use `<default>` as the realm name. The authentication is performed against the application realm. If the naming read operation does not have the **Everyone** special subject

assigned, you must provide the realm that is used by the administrative applications (the registry used in the global security configuration), as well as the appropriate user and password information in that registry for the lookup operation to succeed.

After the lookup operation succeeds, perform another programmatic login by providing the application realm (or *<default>*) and the user and password information for the appropriate user in the registry that is used by the application. This is similar to the case where the login source is prompt. You must authenticate twice, once for the registry used by the global security configuration (for the naming lookup operation) and again for the registry used by the application to access the EJB.

If `com.ibm.CORBA.validateBasicAuth` is set to `false` in the `$WAS_HOME/profiles/$ProfileName/properties/sas.client.props` file then the programmatic login can use *<default>* as the realm name for both the lookup and the EJB operations. The actual authentication occurs only when the resource is accessed on the server side, in which case the realm is calculated based on the resource that is accessed.

The new security domain support for WebSphere Application Version 7.0 does not change the current application security programming model. However, it provides more flexibility and capabilities such as the following:

- User applications can still find the user registry object by using the naming lookup for “UserRegistry”. For the registry object used by administrative applications, the naming lookup for “AdminUserRegistry” can be used.
- The application usage of the JAAS login configuration does not change in a multiple domain setup. However, if an application must refer to the JAAS configuration that is specified at the domain level, the administrator and the deployer of that application must make sure that this domain is configured with the JAAS configurations that are required by the application.
- If an application needs to communicate with other applications using different realms, trust relationship should be established for both inbound and outbound communications when using the LTPA tokens. Read about “Cross realm communication” on page 89 for more information.
- When using programmatic login in the applications, if you want to login to the realm used by the application, use *<default>* as the realm name or provide the realm name that the application is using. If you need to login to the global realm, you must provide the global realm name. If you provide any other realm, only a basic authentication subject is created. When the request actually flows to the server hosting that realm, the actual authentication of the user occurs if that server hosts the realm. If the server does not host the realm, the login fails.

Application deployment in multiple domains configurations

When deploying an application in a multiple domain setup, all of the modules in the application should be installed in the servers or clusters that belong to the same security domain. If not, depending on the security attributes configured in these security domains, inconsistent behavior can result. For example, if the domains contain different user registries, the users and groups information can be different, which can cause inconsistent behavior when accessing the modules. Another example is when the JAAS data is different between the security domains. The JAAS configurations is not accessible from all of the modules in the application. The security runtime code and the command tasks rely on one domain being associated with an application when dealing with attributes such as user registry, JAAS login configurations, J2C authentication data, and authorization.

In most cases, application deployment fails when an application is deployed across different domains. However, since this was possible in earlier releases of WebSphere Application Server when only a few attributes were supported at the server level, the deployment tool first checks for attributes that are configured at the domains. If the attributes in the domain are the same as those supported in previous releases, the administrative console requests confirmation to ensure that you want to deploy application modules across multiple security domains. Unless there is an absolute requirement to deploy the applications across different domains, stop the deployment and select the servers and clusters in the same security domain.

Cross realm communication

When applications communicate using the RMI/IIOP protocol and LTPA is the authentication mechanism, the LTPA token is passed between the servers involved. The LTPA token contains the realm-qualified `accessId`, (also called the *accessId*), of the user who is logging into the front-end application. When this token is received by the downstream server it attempts to decrypt the token. If the LTPA keys are shared between the two servers, decryption succeeds and the `accessId` of the user is obtained from the token. The realm in the `accessId` is checked with the current realm that is used by the application. If the realms match, the LTPA token validation succeeds and it proceeds with the authorization. If the realms do not match, the token validation fails since the user from the foreign realm cannot be validated in the current realm of the application. If applications are not supposed to communicate with each other when using RMI/IIOP and the LTPA authentication mechanism, you do not have to do anything further.

If you do want the cross realm communication to succeed when using RMI/IIOP and LTPA tokens, you must first establish trust between the realms involved, both for inbound and outbound communications.

For the server originating the request, its realm must have the realms that it can trust to send the token to. This is referred to as `outboundTrustedRealms`. For the server receiving the request, its realm needs to trust the realms that it can receive LTPA tokens from. This is referred to as `inboundTrustedRealms`.

Outbound trusted realms can be established using the `addTrustedRealms` command with the `-communicationType` option set to `outbound`. It can also be established in the administrative console by clicking **Trusted authentication realms - outbound** on the **CSiv2 outbound communications** panel.

Inbound trusted realms can be established using the same `addTrustedRealms` command task with the `-communicationType` option set to `inbound`. It can also be established by using the administrative console.

The figure below shows the communication between applications that use different user realms (registries) using RMI/IIOP. In this example, application `app1` (for example, a servlet) is configured to use the `realm1` user registry. The `app2` application (for example, an EJB) is configured to use the `realm2` user registry. The user (`user1`) initially logs into the servlet in `app1`, which then attempts to access an EJB in `app2`. The following must be set:

- In `Domain1`, `realm1` should trust `realm2` for the outbound communication.
- In `Domain2`, `realm2` should trust `realm1` for the inbound communication.
- The `accessId` for `user1` should be configured in the authorization table for `app2`.

When the LTPA token that contains the `accessId` of `user1` is received by `app2`, it decrypts the token. Both of the servers share the same LTPA keys. The LTPA token then ensures that the foreign realm is a trusted realm, and performs the authorization based on the `accessId` of `user1`. If security attribute propagation is not disabled, then the group information of `user1` is also propagated to `app2`. The groups can be used for the authorization check, provided that the authorization table contains the group information. You can associate a special subject, `AllAuthenticatedInTrustedRealms`, to the roles instead of adding individual users and groups to the authorization table.

If the applications in the above example are deployed in different cells, you must do the following:

- Share the LTPA keys between the cells.
- Update the authorization table for `app2` with foreign users and groups `accessIds` by using the **wsadmin** utility. The administrative console does not have access to the realms outside of the scope of the cell.




Figure 4. Cross realm communication in a multiple realm environment

Once trust has been established between the realms, when the server receives the LTPA token and the token is decrypted, it checks to see if the foreign realm is in its inbound trusted realms list. If it is trusted, the authentication succeeds. However, since it is a foreign realm, it does not go search the user registry to gather information about the user. Whatever information is in the LTPA token is used to authorize the user.

The only information in the LTPA token is the unique id of the user. This unique id of the user should exist in the authorization table for this application. If it does, authorization succeeds. However, if attribute propagation is enabled, additional authorization attributes (groups that this user belongs to) for the user are sent from the originating server to the receiving server. These additional attributes are used to make the access decisions. If the groups information exists in the propagation tokens it is used when making the authorization decision.

As previously mentioned, the information about the users and or the groups from the trusted realms should exist in the authorization table of the receiving application. Specifically, the accessId of the users and or groups should exist in the binding file of the application. This must be the case when the application is deployed. In the administrative console, when an application is deployed in a domain you can add the accessIds of the users and groups from any of its trusted realms to the authorization table.

You also have an option to associate a special subject, `AllAuthenticatedInTrustedRealms`, to the roles instead of adding individual users and groups. This is similar to the `AllAuthenticated` special subject that is currently supported. The difference is that the `AllAuthenticated` special subject refers to users in the same realm as the application while the `AllAuthenticatedInTrustedRealms` special subject applies to all of the users in the trusted realms and in the realm of the application.

You can associate the accessId by using the `$AdminApp` install script. Because the accessId takes a unique format, use the command `task listRegistryUsers` with `displayAccessIds` set to `true`. If an invalid name or format is entered in this field, the authorization fails.

User and group information from the trusted realms is obtained by the deployment manager since it has access to all of the user registry configurations in all domains. However, in certain situations it is not possible to obtain the users and group information.

For example, if a server hosted on an external node is using localOS as the registry for its domain, the deployment manager cannot obtain the users and groups information unless it is running in the same operating system setup. The external operating system should be contacted to obtain this information. This can be done by directly invoking the registry in the server associated with that domain. The servers associated with the domain have to be started for this to work. You also must set the property, `com.ibm.websphere.allowRegistryLookupOnProcess`, to `true` in the top-level security custom properties. When this property is set, the deployment manager code searches one of the servers that is associated with the security domain and obtains the users and groups information directly from it. This is possible by calling an MBean in one of the servers.

If the MBean in any of the servers that are using that domain cannot be accessed, the administrative console displays a panel where you can enter the user and `accessId` information manually for each user and group. It is important that the correct `accessId` format be entered in this field. The `accessId` format for the user is `user:realmName/userUniqueId`. The `realmName` is the name of the realm where the user resides, and the `userUniqueId` is the `uniqueId` that represents the user, depending on the registry that is used.

For example, for LDAP, the `uniqueUserId` is the Distinguished Name (DN), for the Windows localOS registry and is the SID of the user. For Unix platforms, it is the UID. For custom registries, it depends on the implementation.

Similarly, for groups, the `accessId` format is `group:realmName/groupUniqueId`. As previously mentioned, use the `listRegistryUsers` and `listRegistryGroups` command with the `-displayAccessIds` option set to `true` so that you can obtain the correct format for the domain or realm that you are interested in.

Once users and groups from the trusted realms or the `AllAuthenticatedInTrustedRealms` special subject is added to the authorization table of the application, it is ready to accept requests from other applications that are using any of its trusted realms. The LTPA token validation on the receiving server first checks to make sure that the realm is trusted. The authorization engine then checks to see if the external user and/or the groups or the `AllAuthenticatedInTrustedRealms` special subject are given access to the roles needed to access the resource. If true, access is granted.

Cross realm communication is only applicable when using the WebSphere built-in authorization. If you are using other authorization engines including SAF for z/OS, any cross realm authorization can be achieved by implementing custom login modules that map external users to users in its own repository.

Federating a node with security domains

When a security domain is configured in the base version and is federated to a cell, the security domain configured at the base version is also configured for that server in the cell. The same domain security configuration can be used by the server before and after the federation. If a base server is to be federated to a cell, the resource assigned to the security domain should be the server scope instead of the cell scope.

If the base server is expected to be registered with an Administrative Agent process, use the cell scope as the resource if the intention is to have all of the servers in the base profile use this security domain.

If during federation the security domain at the base already exists at the cell level, the `addNode` command fails. You can use the `-excludesecuritydomains` option not to include the security domain during federation.

When the federated node is removed from a cell, the resources in that node should be removed from the security domains. If security domains have clusters associated with them that span nodes, the nodes are not removed. You can always remove resources from the security domains or any domains that are not used by using scripting commands or the administrative console.

Security domains in a mixed-version environment

You should create security domains once all of the nodes have been migrated to the latest version. This is especially true if there is a need to associate the cell with a domain. However, if you want to create security domains in a mixed-version environment, be aware of the following:

- If a cell-wide domain is created in a mixed version setup, a domain called *PassThroughToGlobalSecurity* is created automatically. All mixed clusters are assigned to this domain at the time of the creation of the cell-wide domain. This *PassThroughToGlobalSecurity* domain is special in the sense that attributes cannot be added to it, only resources can be assigned to it.

All resources assigned to the *PassThroughToGlobalSecurity* domain use the global security configuration information. Whenever a node in the mixed version setup is migrated to the latest version, the servers and clusters in these nodes are added to this domain. Applications in all of the servers and clusters in these nodes do not use the cell-wide domain; they instead use the global security configuration before and after migration.

If any of these servers need to use the cell-wide domain, you must remove these resources from this *PassThroughToGlobalSecurity* domain. New servers and clusters that are created in the migrated node use the cell-wide domain, not the *PassThroughToGlobalSecurity* domain. As a result, you have a mix of servers and clusters, some of them using global security configuration and some using the cell-wide domain.

- Once a cell-wide domain is created, adding any old version cluster members to a WebSphere Application Server Version 7.0 cluster is restricted since this action makes it a mixed cluster. This restriction also holds true when a WebSphere Application Server Version 7.0 cluster is associated with a domain, and a previous version cluster member is added to this cluster. This restriction is needed to avoid associating a security domain to a mixed cluster.
- If possible, you should create a cell-wide domain after all of the nodes have been migrated. In this case, the cell-wide domain is applicable to the entire cell and not just to parts of it. This also eliminates the need to create the *PassThroughToGlobalSecurity* domain and the mixed cluster scenario with security domains.

Modifying security domains

Use the administrative console tasks or scripting commands to modify security domains. For a complete list of administrative tasks and scripting commands, see the links in "Related tasks" at the bottom of this document.

Once a security domain is created and associated to a set of scopes, the servers associated with this new domain must be restarted. After the restart, the applications in the scopes associated with the new domain use the security attributes defined in the domain.

Changes to any of the domain attributes requires the restart of all of the scopes assigned to it. If new scopes are added they also need to be restarted. Any modifications to the domain configuration, either to the security attributes or to the scopes, has impacts on those applications that are using the domain configuration.

Before you make modifications to an existing domain, consider the following potential impacts. For example, if a user registry that is configured at a domain is removed, and the servers restarted, the user registry from the cell-wide domain (if one is defined), or the global security configuration is then used. This can impact application authentication and authorization. Users and groups associated with an application might no longer be valid in the new registry. Another example to consider is when JAAS configurations are removed from a domain. Applications that rely on this are no longer be able to use the JAAS configurations. Whenever a security configuration is changed it might impact your applications, so all security configuration changes should be made with the utmost care.

Creating new multiple security domains

You can create multiple security domains in your configuration. By creating multiple security domains, you can configure different security attributes for administrative and user applications within a cell environment.

Before you begin

Only users assigned to the administrator role can create new multiple security domains. Enable global security in your environment before creating new multiple security domains.

Read about “Multiple security domains” on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

About this task

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different applications can use different security attributes like user registry or login configurations.

Use multiple security domains to achieve the following goals:

- Configure different security attributes for administrative and user applications within a cell
- Consolidate server configurations by managing different security configurations within a cell
- Restrict access between applications with different user registries, or configure trust relationships between applications to support communication across registries

Perform the following steps to create a new security domain using the administrative console:

1. Click **Security > Security domains**.
2. On the Security domains collection page, click **New**.
3. Specify a unique name for the domain. A domain name must be unique within a cell and cannot contain an invalid character. This field is required.
4. Specify a unique description for the domain. After you click **Apply** you are returned to the Security domains detail page
5. Under Assigned Scopes, assign the security domain to the entire cell or select the specific servers, clusters, and service integration buses to include in the security domain.
6. Customize your security configuration by specifying security attributes for your new domain and by assigning it to cell resources.

You can change security attributes such as the following:

Application Security

Specifies the settings for application security and Java 2 security. You can use the global security settings or customize the settings for a domain.

Select **Enable application security** to enable or disable security this choice for user applications. When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

Java 2 Security

Select **Java 2 security** to enable or disable Java 2 security at the domain level. This choice enables or disables Java 2 security at the process (JVM) level so that all applications (both administrative and user) can enable or disable Java 2 security.

User realm

This section enables you to configure the user registry for the security domain. You can separately configure any registry except the federated registry that is used at the domain level. The federated repository can only be configured at the global level but can be used at the domain level. Read about “Multiple security domains” on page 78 for more information.

Trust association

When you configure the trust association interceptor (TAI) at a domain level, the interceptors configured at the global level are copied to the domain level for convenience. You can modify the interceptor list at the domain level to fit your needs. Only configure those interceptors that are to be used at the domain level.

SPNEGO Web Authentication

The SPNEGO Web authentication, which enables you to configure SPNEGO for Web resource authentication, can be configured at the domain level.

Note: In WebSphere Application Server Version 6.1, a TAI that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

RMI/IOP Security

The RMI/IOP security attribute refers to the CSv2 (Common Secure Interoperability version 2) protocol properties. When you configure these attributes at the domain level, the RMI/IOP security configuration at the global level is copied for convenience.

You can change the attributes that need to be different at the domain level. The Transport layer settings for CSv2 inbound communications should be the same for both the global and the domain levels. If they are different, the domain level attributes are applied to all of the application in the process.

JAAS application logins

Specifies the configuration settings for the Java Authentication and Authorization Service (JAAS) application logins. You can use the global security settings or customize the settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

JAAS system logins

Specifies the configuration settings for the JAAS system logins. You can use the global security settings or customize the configuration settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

JAAS J2C authentication

Specifies the configuration settings for the JAAS J2C authentication data. You can use the global security settings or customize the settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

Authentication Mechanism Attributes

Specifies the various cache settings that need to be applied at the domain level.

Select **Authentication cache settings** to specify your authentication cache settings. The configuration specified on this panel is applied only to this domain.

Select **LTPA Timeout** to configure a different LTPA timeout value at the domain level. The default timeout value is 120 minutes, which is set at the global level. If the LTPA timeout is set at the domain level, any token that is created in the security domain when accessing user applications is created with this expiration time.

When **Use realm-qualified user names** is enabled, user names returned by methods such as `getUserPrincipal()` are qualified with the security realm (user registry) used by applications in the security domain.

Authorization Provider

You can configure an external third party JACC (Java Authorization Contract for Containers) provider at the domain level. Tivoli Access Manager's JACC provider can only be configured at the global level. Security domains can still use it if they do not override the authorization provider with another JACC provider or with the built-in native authorization.

Custom properties

Set custom properties at the domain level that are either new or different from those at the global level. By default, all of the custom properties at the global security configuration can be accessed by all of the applications in the cell. The security runtime code first checks for the custom property at the domain level. If it does not find it, it then attempts to obtain the custom property from the global security configuration.

7. Click **Apply**.
8. After you have saved your configuration changes, restart the server for your changes to take effect.

Deleting multiple security domains

You can delete multiple security domains from your configuration. You must remove the resources assigned to the security domains before deleting them. Only remove those security domains that are not needed in your security configuration.

Before you begin

Only users assigned to the administrator role can delete security domains. Enable global security in your environment before deleting security domains.

Read about "Multiple security domains" on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

About this task

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different applications can use different security attributes like user registry or login configurations.

Perform the following steps to delete an existing security domain using the administrative console:

Note: Only delete the security domains after first removing any resources associated with them. The servers impacted should be restarted.

1. Click **Security > Security domains**.
2. On the Security domains collection page, select a domain to delete.
3. Click **Delete**.

Copying multiple security domains

You can copy selected multiple security domains from the domain collection to create a new domain. This is useful if you want to create a domain that is similar to a previous domain. However, you might want to make a few slight adjustments. When copying an existing domain, you must supply a unique domain name for the new one.

Before you begin

Only users assigned to the administrator role can copy or create new multiple security domains. Enable global security in your environment before copying multiple security domains.

Read about “Multiple security domains” on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

About this task

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different applications can use different security attributes like user registry or login configurations.

Use multiple security domains to achieve the following goals:

- Configure different security attributes for administrative and user applications within a cell
- Consolidate server configurations by managing different security configurations within a cell
- Restrict access between applications with different user registries, or configure trust relationships between applications to support communication across registries

Perform the following steps to copy an existing security domain using the administrative console:

1. Click **Security > Security domains**.
2. Optional: From Preferences, you can select the maximum number of rows to display when the domain collection is large. The default number of rows is 20. Rows that exceed that number appear on subsequent pages.
3. Select a domain to copy.
4. Click **Copy Selected Domain...** to copy an existing domain from the collection. You can optionally select **Copy Global Security..** to copy an existing domain and have it maintain its global security settings (collection selections are ignored). A new domain name is also required if you choose this option.
5. Specify a unique name for the domain. This field is required. A domain name must be unique within a cell and cannot contain an invalid character.
6. Specify a unique description for the domain.
7. Click **Apply**. After you click **Apply** you are returned to the Security domains detail page
8. Under Assigned Scopes, assign the security domain to the entire cell or select the specific servers, clusters, and service integration buses to include in the security domain.

9. Customize your security configuration by specifying security attributes for your new domain and by assigning it to cell resources.

You can change security attributes such as the following:

Application Security

Specifies the settings for application security and Java 2 security. You can use the global security settings or customize the settings for a domain.

Select **Enable application security** to enable or disable security this choice for user applications. When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

Java 2 Security

Select **Java 2 security** to enable or disable Java 2 security at the domain level. This choice enables or disables Java 2 security at the process (JVM) level so that all applications (both administrative and user) can enable or disable Java 2 security.

User realm

This section enables you to configure the user registry for the security domain. You can separately configure any registry except the federated registry that is used at the domain level. The federated repository can only be configured at the global level but can be used at the domain level. Read about “Multiple security domains” on page 78 for more information.

Trust association

When you configure the trust association interceptor (TAI) at a domain level, the interceptors configured at the global level are copied to the domain level for convenience. You can modify the interceptor list at the domain level to fit your needs. Only configure those interceptors that are to be used at the domain level.

SPNEGO Web Authentication

The SPNEGO Web authentication, which enables you to configure SPNEGO for Web resource authentication, can be configured at the domain level.

Note: In WebSphere Application Server Version 6.1, a TAI that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

RMI/IIOP Security

The RMI/IIOP security attribute refers to the CSv2 (Common Secure Interoperability version 2) protocol properties. When you configure these attributes at the domain level, the RMI/IIOP security configuration at the global level is copied for convenience.

You can change the attributes that need to be different at the domain level. The Transport layer settings for CSv2 inbound communications should be the same for both the global and the domain levels. If they are different, the domain level attributes are applied to all of the application in the process.

JAAS application logins

Specifies the configuration settings for the Java Authentication and Authorization Service (JAAS) application logins. You can use the global security settings or customize the settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

JAAS system logins

Specifies the configuration settings for the JAAS system logins. You can use the global security settings or customize the configuration settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

JAAS J2C authentication

Specifies the configuration settings for the JAAS J2C authentication data. You can use the global security settings or customize the settings for a domain.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

Authentication Mechanism Attributes

Specifies the various cache settings that need to be applied at the domain level.

Select **Authentication cache settings** to specify your authentication cache settings. The configuration specified on this panel is applied only to this domain.

Select **LTPA Timeout** to configure a different LTPA timeout value at the domain level. The default timeout value is 120 minutes, which is set at the global level. If the LTPA timeout is set at the domain level, any token that is created in the security domain when accessing user applications is created with this expiration time.

When **Use realm-qualified user names** is enabled, user names returned by methods such as `getUserPrincipal()` are qualified with the security realm (user registry) used by applications in the security domain.

Authorization Provider

You can configure an external third party JACC (Java Authorization Contract for Containers) provider at the domain level. Tivoli Access Manager's JACC provider can only be configured at the global level. Security domains can still use it if they do not override the authorization provider with another JACC provider or with the built-in native authorization.

Custom properties

Set custom properties at the domain level that are either new or different from those at the global level. By default, all of the custom properties at the global security configuration can be accessed by all of the applications in the cell. The security runtime code first checks for the custom property at the domain level. If it does not find it, it then attempts to obtain the custom property from the global security configuration.

10. Click **Apply**.

11. After you have saved your configuration changes, restart the server for your changes to take effect.

Configuring inbound trusted realms for multiple security domains

You can configure which realms to grant inbound trust to for multiple security domains. The trust relationship between realms is used when communicating with Lightweight Third-Party Authentication (LTPA) tokens. Once a LTPA token is decrypted by the receiving server, the realm in the token is checked to see if it is trusted. If it is not, the validation of the token fails. A *realm* represents a user registry in WebSphere Application Server.

Before you begin

For information on cross realm communications, read the section in “Multiple security domains” on page 78.

Only users assigned to the administrator role can configure multiple security domains. Enable global security in your environment before configuring multiple security domains.

Perform the following steps to grant inbound trusted realms for multiple security domains using the administrative console:

1. Click **Security > Security domains**.
2. Select a domain to edit or create a new one. Under Security Attributes, click **User realm**.
3. Click **Customize for this domain**.
4. Under Related Items, select **Trusted authentication realms - inbound**.
5. Select **Trust all realms (including those external to this cell)** or **Trust realms as indicated below**. If Kerberos authentication is enabled, and you have cross realms or trusted realms, you must add the Kerberos trusted realm by selecting **Trust realms as indicated below**.
6. Click **Apply**.

What to do next

The realms you selected to trust accept messages from other trusted realms but do not accept messages from untrusted realms. Select **Add External Realm** to add trust for realms that are external to this cell.

Configure security domains

Use this page to configure the security attributes of a domain and to assign the domain to cell resources. For each security attribute, you can use the global security settings or customize settings for the domain.

To view this administrative console page, click **Security > Security domains**. On the Security domains collection page, select an existing domain to configure, create a new one, or copy an existing domain.

Read about “Multiple security domains” on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

Related concepts

“Multiple security domains” on page 78

The WebSphere Security Domains (WSD) provide the flexibility to use different security configurations in WebSphere Application Server. The WSD is also referred to as multiple security domains, or simply, security domains. You can configure different security attributes, such as the UserRegistry, for different applications.

Related reference

“Standalone LDAP registry settings” on page 122

Use this page to configure Lightweight Directory Access Protocol (LDAP) settings when users and groups reside in an external LDAP directory.

“Configuration entry settings for Java Authentication and Authorization Service” on page 408

Use this page to specify a list of Java Authentication and Authorization Service (JAAS) login configurations for the application code to use, including Java 2 Platform, Enterprise Edition (J2EE) components such as enterprise beans, JavaServer Pages (JSP) files, servlets, resource adapters, and message-driven beans (MDBs).

“Java 2 Connector authentication data entry settings” on page 420

Use this page as a central place for administrators to define authentication data, which includes user identities and passwords. These values can reference authentication data entries by resource adapters, data sources, and other configurations that require authentication data using an alias.

“External Java Authorization Contract for Containers provider settings” on page 539

Use this page to configure the application server to use an external Java Authorization Contract for Containers (JACC) provider. For example, the policy class name and the policy configuration factory class name are required by the JACC specification.

“Security domains collection” on page 107

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different application servers can use different security attributes like user registry or login configurations.

Related information

“Security custom properties” on page 62

Use this page to understand the predefined custom properties that are related to security.

Name

Specifies a unique name for the domain. This name can not be edited after the initial submission.

A domain name must be unique within a cell and cannot contain an invalid character.

Description

Specifies a description for the domain.

Assigned Scopes

Select to display the cell topology. You can assign the security domain to the entire cell or select the specific clusters, nodes and service integration buses to include in the security domain.

If you select **All scopes**, the entire cell topology is displayed.

If you select **Assigned scopes**, the cell topology is displayed with those servers and clusters that are assigned to the current domain.

The name of an explicitly assigned domain appears next to any resource. Checked boxes indicate resources that are currently assigned to the domain. You also can select other resources and click **Apply** or **OK** to assign them to the current domain.

A resource that is not checked (disabled) indicates that it is not assigned to the current domain and must be removed from another domain before it can be enabled for the current domain.

If a resource does not have an explicitly-assigned domain, it uses the domain assigned to the cell. If no domain is assigned to the cell, then the resource uses global settings.

Cluster members cannot be individually assigned to domains; the entire cluster uses the same domain.

Application Security:

Select **Enable application security** to enable or disable security for user applications. You can use the global security settings or customize the settings for a domain.

When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

Enable application security

Enables security for the applications in your environment. This type of security provides application isolation and requirements for authenticating application users

In previous releases of WebSphere Application Server, when a user enabled global security, both administrative and application security were enabled. In WebSphere Application Server Version 6.1, the previous notion of global security were split into administrative security and application security, each of which you can enable separately.

As a result of this split, WebSphere Application Server clients must know whether application security is disabled at the target server. Administrative security is enabled, by default. Application security is disabled, by default. To enable application security, you must enable administrative security. Application security is in effect only when administrative security is enabled.

When this selection is disabled, all of the EJBs and Web applications in the security domain are no longer protected. Access is granted to these resources without user authentication. When you enable this selection, the J2EE security is enforced for all of the EJBs and Web applications in the security domain. The J2EE security is only enforced when Global Security is enabled in the global security configuration, (that is, you cannot enable application security without first enabling Global Security at the global level).

Java 2 security:

Select **Use Java 2 security** to enable or disable Java 2 security at the domain level or to assign or add properties related to Java 2 security. You can use the global security settings or customize the settings for a domain.

This choice enables or disables Java 2 security at the process (JVM) level so that all applications (both administrative and user) can enable or disable Java 2 security.

Use global security settings

Select to specify the global security settings that are being used.

Customize for this domain

Select to specify the settings that are defined in the domain, such as options to enable application and Java 2 security and to use realm-qualified authentication data.

Use Java 2 security to restrict application access to local resources

Select to specify whether to enable or disable Java 2 security permission checking. By default, access to local resources is not restricted. You can choose to disable Java 2 security, even when application security is enabled.

When the **Use Java 2 security to restrict application access to local resources** option is enabled and if an application requires more Java 2 security permissions than are granted in the default policy, the application might fail to run properly until the required permissions are granted in either the app.policy file or the was.policy file of the application. AccessControl exceptions are generated by applications that do not have all the required permissions.

Warn if applications are granted custom permissions

Specifies that during application deployment and application start, the security runtime issues a warning if applications are granted any custom permissions. Custom permissions are permissions that are defined by the user applications, not Java API permissions. Java API permissions are permissions in the java.* and javax.* packages.

The application server provides support for policy file management. A number of policy files are available in this product, some of them are static and some of them are dynamic. Dynamic policy is a template of permissions for a particular type of resource. No code base is defined and no relative code base is used in the dynamic policy template. The real code base is dynamically created from the configuration and run-time data. The filter.policy file contains a list of permissions that you do not want an application to have according to the J2EE 1.4 specification.

Note: You cannot enable this option without enabling the **Use Java 2 security to restrict application access to local resources** option.

Restrict access to resource authentication data

This option is disabled if Java 2 security has not been enabled.

Consider enabling this option when both of the following conditions are true:

- Java 2 security is enforced.
- The application code is granted the accessRuntimeClasses WebSphereRuntimePermission permission in the was.policy file found within the application enterprise archive (EAR) file. For example, the application code is granted the permission when the following line is found in your was.policy file:

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

The **Restrict access to resource authentication data** option adds fine-grained Java 2 security permission checking to the default principal mapping of the WSPrincipalMappingLoginModule implementation. You must grant explicit permission to Java 2 Platform, Enterprise Edition (J2EE) applications that use the WSPrincipalMappingLoginModule implementation directly in the Java Authentication and Authorization Service (JAAS) login when **Use Java 2 security to restrict application access to local resources** and the **Restrict access to resource authentication data** options are enabled.

Default: Disabled

User Realm:

This section enables you to configure the user registry for the security domain. You can separately configure any registry except the federated registry that is used at the domain level. The federated repository can only be configured at the global level but can be used at the domain level.

When configuring a registry at the domain level you can choose to define your own realm name for the registry. The realm name distinguishes one user registry from another. The realm name is used in multiple places – in the Java client login panel to prompt the user, in the authentication cache, and when using native authorization.

At the global configuration level, the system creates the realm for the user registry. In previous releases of WebSphere Application Server, only one user registry is configured in the system. When you have multiple security domains you can configure multiple registries in the system. For the realms to be unique in these domains, configure your own realm name for a security domain. You also can choose the system to create a unique realm name if it is certain to be unique. In the latter case, the realm name is based on the registry that is being used.

Trust Association:

Select to specify the settings for the trust association. Trust association is used to connect reversed proxy servers to the application servers.

Trust association enables the integration of IBM WebSphere Application Server security and third-party security servers. More specifically, a reverse proxy server can act as a front-end authentication server while the product applies its own authorization policy onto the resulting credentials that are passed by the proxy server.

Tivoli Access Manager's trust association interceptors can only be configured at the global level. The domain configuration can also use them, but cannot have a different version of the trust association interceptor. Only one instance of Tivoli Access Manager's trust association interceptors can exist in the system.

Note: The use of trust association interceptors (TAIs) for Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) authentication is deprecated. The SPNEGO web authentication panels provide a much easier way to configure SPNEGO.

Interceptors

Select to access or to specify the trust information for reverse proxy servers.

Enable trust association

Select to enable the integration of IBM WebSphere Application Server security and third-party security servers. More specifically, a reverse proxy server can act as a front-end authentication server while the product applies its own authorization policy onto the resulting credentials that are passed by the proxy server.

SPNEGO Web Authentication:

Specifies the settings for Simple and Protected GSS-API Negotiation (SPNEGO) as the Web authentication mechanism.

The SPNEGO Web authentication, which enables you to configure SPNEGO for Web resource authentication, can be configured at the domain level.

Note: In WebSphere Application Server Version 6.1, a TAI that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

RMI/IIOP Security:

Specifies the settings for Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP).

An Object Request Broker (ORB) manages the interaction between clients and servers, using the Internet InterORB Protocol (IIOP). It enables clients to make requests and receive responses from servers in a network-distributed environment.

When you configure these attributes at the domain level, the RMI/IIOP security configuration at the global level is copied for convenience. You can change the attributes that need to be different at the domain level. The Transport layer settings for CSiv2 inbound communications should be the same for both the global and the domain levels. If they are different, the domain level attributes are applied to all of the applications in the process.

When a process communicates with another process with a different realm, the LTPA authentication and the propagation tokens are propagated to the downstream server unless that server is listed in the outbound trusted realms list. This can be done using the **Trusted authentication realms – outbound** link on the **CSiv2 outbound communication** panel.

CSiv2 inbound communications

Select to specify authentication settings for requests that are received and transport settings for connections that are accepted by this server using the Object Management Group (OMG) Common Secure Interoperability (CSI) authentication protocol.

WebSphere Application Server enables you to specify Internet Inter-ORB Protocol (IIOP) authentication for both inbound and outbound authentication requests. For inbound requests, you can specify the type of accepted authentication, such as basic authentication.

CSiv2 outbound communications

Select to specify authentication settings for requests that are sent and transport settings for connections that are initiated by the server using the Object Management Group (OMG) Common Secure Interoperability (CSI) authentication protocol.

WebSphere Application Server enables you to specify Internet Inter-ORB Protocol (IIOP) authentication for both inbound and outbound authentication requests. For outbound requests, you can specify properties such as type of authentication, identity assertion or login configurations that are used for requests to downstream servers.

JAAS Application logins

Select to define login configurations that are used by JAAS.

The JAAS application logins, the JAAS system logins, and the JAAS J2C authentication data aliases can all be configured at the domain level. By default, all of the applications in the system have access to the JAAS logins configured at the global level. The security runtime first checks for the JAAS logins at the domain level. If it does not find them, it then checks for them in the global security configuration. Configure any of these JAAS logins at a domain only when you need to specify a login that is used exclusively by the applications in the security domain.

For JAAS and custom properties only, once global attributes are customized for a domain they can still be used by user applications.

Do not remove the ClientContainer, DefaultPrincipalMapping, and WSLogin login configurations because other applications might use them. If these configurations are removed, other applications might fail.

Use global and domain-specific logins

Select to specify the settings that are defined in the domain, such as options to enable application and Java 2 security and to use realm-qualified authentication data.

JAAS System Logins:

Specifies the configuration settings for the JAAS system logins. You can use the global security settings or customize the configuration settings for a domain.

System Logins

Select to define the JAAS login configurations that are used by system resources, including the authentication mechanism, principal mapping, and credential mapping

JAAS J2C Authentication Data:

Specifies the settings for the JAAS J2C authentication data. You can use the global security settings or customize the settings for a domain.

Java 2 Platform, Enterprise Edition (J2EE) Connector authentication data entries are used by resource adapters and Java DataBase Connectivity (JDBC) data sources.

Use global and domain-specific entries

Select to specify the settings that are defined in the domain, such as options to enable application and Java 2 security and to use realm-qualified authentication data.

Authentication Mechanism Attributes:

Specifies the various cache settings that must be applied at the domain level.

- Authentication cache settings - use to specify your authentication cache settings. The configuration specified on this panel is applied only to this domain.
- LTPA Timeout - You can configure a different LTPA timeout value at the domain level. The default timeout value is 120 minutes, which is set at the global level. If the LTPA timeout is set at the domain level, any token that is created in the security domain when accessing user applications is created with this expiration time.
- Use realm-qualified user names - When this selection is enabled, user names returned by methods such as `getUserPrincipal()` are qualified with the security realm (user registry) used by applications in the security domain.

Authorization Provider:

Specifies the settings for the authorization provider. You can use the global security settings or customize the settings for a domain.

You can configure an external third party JACC (Java Authorization Contract for Containers) provider at the domain level. Tivoli Access Manager's JACC provider can only be configured at the global level. Security domains can still use it if they do not override the authorization provider with another JACC provider or with the built-in native authorization.

Select either the **Default authorization** or **External authorization using a JAAC provider**. The **Configure** button is only enabled when **External authorization using a JAAC provider** is selected.

Custom properties

Select to specify name-value pairs of data, where the name is a property key and the value is a string.

Set custom properties at the domain level that are either new or different from those at the global level. By default, all of the custom properties at the global security configuration can be accessed by all of the applications in the system. The security runtime code first checks for the custom property at the domain level. If it does not find it, it then attempts to obtain the custom property from the global security configuration.

Web Services Bindings

Click **Default policy set bindings** to set the domain default provider and client bindings.

External realm name

Use this page to add a WebSphere Application Server realm that is external to this cell. The realm is initially not trusted. Use the Trusted authentication realms - inbound page to establish trust.

To view this administrative console page, click **Security > Security domains**. Select a domain to edit or create a new one. Under Security Attributes, click **User realm**. Click **Customize for this domain** and then select a **Realm type**. Click **Configure**. Under Related items, click **Trusted authentication realms - inbound** or **Trusted authentication realms - outbound**. Click **Add External Realm...**

Related reference

“Trust all realms”

Use this page to configure which realms to grant inbound or outbound trust to.

External realm name

Use to specify the name of the realm that is external to the list of realms that are available to receive trust.

Trust all realms

Use this page to configure which realms to grant inbound or outbound trust to.

The inbound trust is required to validate LTPA tokens that contain a foreign realm. The outbound trust is required to send the credential tokens to the trusted realms. For example, if an application using realmA needs to communicate using LTPA with an application using realmB, realmA should have realmB in its outbound trust list and realmB should have realmA in its inbound trust list.

To view this administrative console page, click **Security > Security domains**. Select a domain to edit or create a new one. Under Security Attributes, click **User realm**. Click **Customize for this domain**. Select a realm type and then click **Configure**.

Under Related items, click **Trusted authentication realms - inbound** or **Trusted authentication realms - outbound**.

Trust all realms (including those external to this cell)

Select to trust all of the realms listed on this page, including those external to the cell.

Trust realms as indicated below

Select to trust only those realms that you have selected from the list of realms that are available to receive inbound trust.

Add External Realm...

Select to add realms that are external to this cell to the list of realms that are available to receive inbound trust. When an external realm is added, it is trusted by default. If it is not trusted it is removed from the list.

Security domains collection

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different application servers can use different security attributes like user registry or login configurations.

To view this administrative console page, click **Security > Security domains**.

Read about “Multiple security domains” on page 78 for a better understanding of what multiple security domains are and how they are supported in this version of WebSphere Application Server.

Related reference

“Configure security domains” on page 99

Use this page to configure the security attributes of a domain and to assign the domain to cell resources. For each security attribute, you can use the global security settings or customize settings for the domain.

Related information

“Security custom properties” on page 62

Use this page to understand the predefined custom properties that are related to security.

Maximum rows

Specifies the maximum number of rows that display when the collection is large. The rows that are not displayed appear on the next page.

The default is 20. Rows that exceed the maximum number display on subsequent pages.

Retain filter criteria

Specifies whether to use the same filter criteria entered in the show filter function to display this page the next time you visit it.

Copy selected domain

Select to copy a selected domain from the collection (a new name is required)

Copy global security

Select to create a domain with a copy of the global security settings (collection selections are ignored). A domain name is required.

Authentication cache settings

Use this page to specify your authentication cache settings.

To view this administrative console page, click **Security > Global security > Authentication cache settings**.

Related reference

“Security domains collection”

Security domains provide a mechanism to use different security settings for administrative applications and user applications. They also provide the ability to support multiple security settings so different application servers can use different security attributes like user registry or login configurations.

Enable authentication cache

Specifies whether to disable the authentication cache.

Leave the authentication cache enabled for performance reasons. However, you can disable the authentication cache for debug or measurement purposes. When this choice is disabled, the performance

is impacted since whenever a user is authenticated the user registry is accessed to gather information about the user. New tokens are then created for the user.

Default: Enabled

Cache timeout:

Specifies the time period at which the authenticated credential in the cache expires. Verify that this time period is less than the value for the **Timeout value for forwarded credentials between servers** field (the LTPA timeout).

If the application server infrastructure security is enabled, the security cache timeout can influence performance. The timeout setting specifies how often to refresh the security-related caches. Security information pertaining to beans, permissions, and credentials is cached. When the cache timeout expires, all cached information not accessed within the timeout period is purged from the cache. Subsequent requests for the information result in a database lookup. On occasion, acquiring the information requires invoking a Lightweight Directory Access Protocol (LDAP)-bind or native authentication. Both invocations are relatively costly operations for performance. Determine the best trade-off for the application by looking at usage patterns and security needs for the site.

The default security cache timeout value is 10 minutes. If you have a small number of users, it should be set higher than that, or lower if a larger number of users.

The LTPA timeout value should not be set lower than the security cache timeout. The LTPA timeout value should be set higher than the orb request timeout value. However, there is no relation between the security cache timeout value and the orb request timeout value.

Default: 10 minutes

Initial cache size:

Specifies the initial size of the hash table caches.

A higher number of available hash values might decrease the occurrence of hash collisions. A hash collision results in a linear search for the hash bucket, which might decrease the retrieval time. If several entries compose a hash table cache, create a table with a larger capacity that supports more efficient hash entries instead of allowing automatic rehashing determine the growth of the table. Rehashing causes every entry to move each time.

Default: 50

Maximum cache size

Indicates the maximum size of the cache.

After this limit is reached, the least used entries are removed from the cache to make space for the new entries.

Default: 25000

Use basic authentication cache keys (password one-way hashed):

Caches the userName and the one-way hashed password as the key lookup in the cache.

Disable this only if you do want this information to be stored in the cache. If this is disabled, every time a user logs in with userName and password, the user registry is accessed, which impacts performance.

Default: True

Use custom cache keys:

Enables custom cache keys to be used as the key lookups in the authentication cache.

Default: True

Chapter 5. Authenticating users

The process of authenticating users involves a user registry and an authentication mechanism. Optionally, you can define trust between WebSphere Application Server and a proxy server, configure single sign-on capability, and specify how to propagate security attributes between application servers.

About this task

The following security topics are covered in this section:

User registries

For information on local operating system, Lightweight Directory Access Protocol (LDAP), custom user registries, and user repositories such as virtual member manager, see “Selecting a registry or repository.”

Trust associations

For more information on trust associations, see “Trust associations” on page 289.

Single sign-on

For more information on single sign-on, see “Single sign-on using LTPA cookies” on page 295.

Security attribute propagation

For more information on propagation tokens, authorization tokens, single sign-on tokens, and authentication tokens, see “Security attribute propagation” on page 448.

The following information is covered in this section:

- Configure a user registry. For more information, see “Selecting a registry or repository.”
- Configure WebSEAL or a custom trust association interceptor. For more information see, “Integrating third-party HTTP reverse proxy servers” on page 293.
- Configure single sign-on. For more information, see “Implementing single sign-on to minimize Web user authentications” on page 297.
- Propagate security attributes. For more information, see “Propagating security attributes among application servers” on page 453.
- Configure the authentication cache. For more information, see “Configuring the authentication cache” on page 466.

What to do next

After completing the configuring the authentication process, you must authorize access to resources. For more information, see Chapter 6, “Authorizing access to resources,” on page 509.

Selecting a registry or repository

Information about users and groups reside in a user registry. In WebSphere Application Server, a user registry authenticates a user and retrieves information about users and groups to perform security-related functions, including authentication and authorization.

Before you begin

Note: During profile creation, either during installation or post-installation, administrative security is enabled by default. The file-based federated user repository is configured as the active user registry. Decide if you want a different user registry.

Before configuring the user registry or repository, decide which user registry or repository to use. You can configure one Active default registry for the Cell.

About this task

WebSphere Application Server provides implementations that support multiple types of registries and repositories including the local operating system registry, a standalone Lightweight Directory Access Protocol (LDAP) registry, a standalone custom registry, and federated repositories.

With WebSphere Application Server, a user registry or a repository, such as a federated repository, authenticates a user and retrieves information about users and groups to perform security-related functions including authentication and authorization.

With WebSphere Application Server, a user registry or repository is used for:

- Authenticating a user using basic authentication, identity assertion, or client certificates
- Retrieving information about users and groups to perform security-related administrative functions, such as mapping users and groups to security roles

In addition to local operating system, LDAP, and Federated repository registries, WebSphere Application Server also provides a plug-in to support any registry by using the custom registry feature. The custom registry feature enables you to configure any user registry that is not made available through the security configuration panels of the WebSphere Application Server.

Configuring the correct registry or repository is a prerequisite to assigning users and groups to roles for applications. When a user registry or repository is not configured, the local operating system registry is used by default. If your choice of user registry is not the local operating system registry, you need to first configure the registry or repository, which is normally done as part of enabling security, restart the servers, and then assign users and groups to roles for all your applications.

WebSphere Application Server supports the following types of user registries:

- Federated repository
- Local operating system
- Standalone Lightweight Directory Access Protocol (LDAP) registry
- Standalone custom registry

The UserRegistry interface is used to implement both the custom registry and the federated repository options for the user account repository. The interface is very helpful in situations where the current user and group information exists in some other formats, for example, a database, and cannot move to local operating system or LDAP registries. In such a case, you can implement the UserRegistry interface so that WebSphere Application Server can use the existing registry for all the security-related operations. The process of implementing a custom registry is a software implementation effort, and it is expected that the implementation does not depend on WebSphere Application Server resource management for its operation. For example, you cannot use an Application Server data source configuration; generally you must invoke database connections and dictate their behavior directly in your code.

Note: WebSphere Application Server has implemented a user registry proxy by using the UserRegistry interface. However, the return values are little different from the interface. For example, `getUniqueId` returns the uniqueID with the realm name wrapped. You cannot use the return value to pass to `getUserSecurityName`, as shown in the following example:

```
// Retrieves the default InitialContext for this server.
javax.naming.InitialContext ctx = new javax.naming.InitialContext();

// Retrieves the local UserRegistry object.
com.ibm.websphere.security.UserRegistry reg =
    (com.ibm.websphere.security.UserRegistry) ctx.lookup("UserRegistry");

// Retrieves the registry uniqueID based on the userName that is specified
// in the NameCallback.
String uniqueid = reg.getUniqueId(userName);
```

```
// Strip the realm name and get real uniqueID
String uid = com.ibm.wsspi.security.token.WSSecurityPropagationHelper.getUserFromUniqueID (uniqueID);

// Retrieves the security name from the user registry based on the uniqueID.
String securityName = reg.getUserSecurityName(uid);
```

You can use a Service Provider Interface (SPI) for this parsing function.

After the applications are assigned users and groups and you need to change the user registries, delete all the users and groups, including any RunAs role, from the applications, and reassign them after changing the registry through the administrative console or by using wsadmin scripting. The following **wsadmin** command, which uses Jacl, removes all of the users and groups from any application:

```
$AdminApp deleteUserAndGroupEntries yourAppName
```

where *yourAppName* is the name of the application. Backing up the old application is advised before performing this operation. However, if both of the following conditions are true, you might be able to switch the registries without having to delete the users and groups information:

- All of the user and group names, including the password for the RunAs role users, in all of the applications match in both user registries.
- The application bindings file does not contain the access IDs which are unique for each user registry even for the same user or group name.

By default, an application does not contain access IDs in the bindings file. These IDs are generated when the applications start. However, if you migrated an existing application from an earlier release, or if you used the wsadmin script to add access IDs for the applications to improve performance, you have to remove the existing user and group information and add the information after configuring the new user registry.

For more information on updating access IDs, see updateAccess IDs in the AdminApp object for scripted administration article.

Note: WebSphere Application Server supports a variety of user registries and repositories on different operating systems. During the user authentication process, you might use non-alphanumeric characters in your user name or password. Restrictions on the use of these non-alphanumeric characters depends on both the underlying operating system and the user registry type. For more information on which non-alphanumeric characters are not supported, see your operating system and user registry or repository documentation.

AIX For example, the following characters are not supported in a user name value:

- `
- #
- =
- \
- :
- "
- ,
- /
- ?
- '
 - A space character

For a comprehensive list of the non-alphanumeric characters that are not supported, see the IBM AIX operating system documentation.

For example, the following characters are not supported in a user name value:

- `
- :
- "
- /
- A space character

Complete one of the following steps to configure your user registry:

- “Configuring local operating system registries” on page 115
- “Configuring Lightweight Directory Access Protocol user registries” on page 119
- “Configuring standalone custom registries” on page 142.
- “Managing the realm in a federated repository configuration” on page 170

What to do next

1. If you are enabling security, make sure that you complete the remaining steps. Verify that the User account repository on the Global security panel is set to the appropriate registry or repository. As the final step, validate the user ID and the password by clicking **Apply** on the Global security panel. Save, stop and start all WebSphere Application Servers.
2. For any changes in user registry panels to be effective, you must validate the changes by clicking **Apply** on the Global security panel. After validation, save the configuration and stop and start all WebSphere Application Servers, including the cells, nodes and all of the application servers. To avoid inconsistencies between the WebSphere Application Server processes, make sure that any changes to the registry or repository are done when all of the processes are running. If any of the processes are down, force synchronization to make sure that the process can start later.

If the server or servers start without any problems, the setup is correct.

Standalone custom registries

A *standalone custom registry* is a customer-implemented registry that implements the UserRegistry Java interface, as provided by the product. A custom-implemented registry can support virtually any type of an account repository from a relational database, flat file, and so on. The custom user registry provides considerable flexibility in adapting product security to various environments where some form of a registry or repository other than federated repositories, standalone Lightweight Directory Access Protocol (LDAP) registry or local operating system registry already exists in the operational environment.

WebSphere Application Server security provides an implementation that uses various local operating system-based registries and various standalone Lightweight Directory Access Protocol (LDAP)-based registries. However, situations can exist where your user and group data resides in other repositories or custom user registries, such as a database, and moving this information to either a local operating system registry or a standalone LDAP registry implementation might not be feasible. For these situations, WebSphere Application Server security provides a service provider interface (SPI) that you can implement to interact with your current registry. The custom registry feature supports any user registry that is not implemented by WebSphere Application Server.

The SPI is the UserRegistry interface. The *UserRegistry interface* is a collection of methods that are required for authorization purposes. These methods authenticate individual users using either a password or certificates and collect information about the user, which are called privilege attributes. This interface also includes methods that obtain user and group information so that they can be given access to resources. When implementing the methods in the interface, you must decide how to map the information that is manipulated by the UserRegistry interface to the information in your registry.

This interface has a set of methods to implement for the product security to interact with your registries for all security-related tasks. The local operating system and LDAP registry implementations that are provided

also implement this interface. Custom user registries are sometimes called the *pluggable user registries* or *custom registries* for short. Your custom user registry implementation is expected to be thread-safe.

Building a custom registry is a software implementation effort. The implementation does not depend on other WebSphere Application Server resources, for example, data sources, for its operation.

Make sure that your implementation of the custom registry does not depend on any WebSphere Application Server components such as data sources, enterprise beans, and so on. Do not have this dependency because security is initialized and enabled prior to most of the other WebSphere Application Server components during startup. If your previous implementation used these components, make a change that eliminates the dependency.

The methods in the UserRegistry interface operate on the following information for users:

User security name

The user name is similar to the user name in the local operating system registries.

This name is used to log in when prompted by a secured application. By default, the Enterprise JavaBeans (EJB) `getCallerPrincipal` method and the `getRemoteUser` and `getUserPrincipal` servlet methods return this name. The user security name is also referred to as *userSecurityName*, *userName*, or *user name*.

Unique user ID

This ID represents a unique identifier for the user, which is required by the UserRegistry interface. The unique ID is similar to the system ID (SID) in Windows systems, the Unique ID (UID) in Linux® and UNIX systems, and the distinguished name (DN) in Lightweight Directory Authentication Protocol (LDAP). This ID is also referred to as *uniqueUserId*. The unique ID is used to make the authorization decisions for protected resources.

Display user name

This name is an optional string that describes a user. The implementation can use display names for informational purposes only; these names are not required to exist or to be unique. The user interface can use the display name to present more information about the user.

Group security name

This name, which represents the security group, is also referred to as *groupSecurityName*, *groupName*, and *group name*.

Unique group ID

The unique ID is the identifier for a group. This name is also referred to as *uniqueGroupId* ID.

Display group name

The display name is an optional string that describes a group.

The topic on UserRegistry interface describes each of the methods in the interface that need implementing. An explanation of each of the methods and their usage in the sample and any changes from the Version 4 interface are provided. The Related references section provides links to all other custom user registries documentation, including a file-based registry sample. The Sample provided is very simple and is intended to familiarize you with this feature. Do not use this sample in an actual production environment.

Configuring local operating system registries

Use these steps to configure local operating system registries.

Before you begin

For detailed information about using the local operating system user registry, see “Local operating system registries” on page 238. These steps set up security based on the local operating system user registry on which WebSphere Application Server is installed.

For security purposes, the WebSphere Application Server provides and supports the implementation for Windows operating system registries, AIX, Solaris and multiple versions of Linux operating systems. The

respective operating system application programming interface (API) are called by the product processes (servers) for authenticating a user and other security-related tasks (for example, getting user or group information). Access to these APIs are restricted to users who have special privileges. These privileges depend on the operating system and are described below.

In WebSphere Application Server Version 6.1, you can use an internally-generated server ID because the Security WebSphere Common Configuration Model (WCCM) model contains a new tag, `internalServerId`. You do not need to specify a server user ID and a password during security configuration except in a mixed-cell environment. See “Administrative roles and naming service authorization” on page 510 for more detailed information about the new internal server ID.

Windows Consider the following issues:

- The server ID needs to be different from the Windows machine name where the product is installed. For example, if the Windows machine name is `vicky` and the security server ID is `vicky`, the Windows system fails when getting the information (group information, for example) for user `vicky`.
- WebSphere Application Server dynamically determines whether the machine is a member of a Windows system domain.
- WebSphere Application Server does not support Windows trusted domains.
- If a machine is a member of a Windows domain, both the domain user registry and the local user registry of the machine participate in authentication and security role mapping.
- The domain user registry takes precedence over the local user registry of the machine and can have undesirable implications if users with the same password exist in both user registries.
- The user that the product processes run under requires the Administrative and Act as part of the operating system privileges to call the Windows operating system APIs that authenticate or collect user and group information. The process needs special authority, which is given by these privileges. The user in this example might not be the same as the security server ID (the requirement for which is a valid user in the registry). This user logs into the machine (if using the command line to start the product process) or the Log On User setting in the services panel if the product processes have started using the services. If the machine is also part of a domain, this user is a part of the Domain Admin group in the domain to call the operating system APIs in the domain in addition to having the Act as part of operating system privilege in the local machine.

Consider the following points:

- **AIX** **UNIX** **Solaris** The user that the product processes run under requires the root privilege. This privilege is needed to call the operating system APIs to authenticate or to collect user and group information. The process needs special authority, which is given by the root privilege. This user might not be the same as the security server ID (the requirement is that it should be a valid user in the registry). This user logs into the machine and is running the product processes.
- **UNIX** The user that enables administrative security must have the root privilege if you use the local operating system registry. Otherwise, a failed validation error is displayed.
- **Linux** You might need to have the password shadow file in your system.

About this task

The following steps are needed to perform this task initially when setting up security for the first time.

1. Click **Security > Global security**.
2. Under User account repository, select **Local operating system** and click **Configure**.
3. Enter a valid user name in the **Primary administrative user name** field. This value is the name of a user with administrative privileges that is defined in the registry. This user name is used to access the administrative console or used by `wsadmin`.
4. Click **Apply**.
5. Select either the **Automatically generated server identity** or **Server identity that is stored in the repository** option. If you select the **Server identity that is stored in the repository** option, enter the following information:

Server user ID or administrative user on a Version 6.0.x node

Specify the short name of the account that is chosen in the second step.

Server user password

Specify the password of the account that is chosen in the second step.

6. Click **OK**.

The administrative console does not validate the user ID and password when you click **OK**. Validation is only done when you click **OK** or **Apply** in the Global security panel. First, make sure that you select **Local operating system** as the available realm definition in the User account repository section, and click **Set as current**. If security was already enabled and you had changed either the user or the password information in this panel, make sure to go to the Global security panel and click **OK** or **Apply** to validate your changes. If your changes are not validated, the server might not start.

Note: Until you authorize other users to perform administrative functions, you can only access the administrative console with the server user ID and password that you specified. For more information, see “Authorizing access to administrative roles” on page 561.

Results

For any changes in this panel to be effective, you need to save, stop, and start all the product servers, including deployment managers, nodes and application servers. If the server comes up without any problems, the setup is correct.

After completed these steps, you have configured WebSphere Application Server to use the local operating system registry to identify authorized users.

What to do next

Complete any remaining steps for enabling security. For more information, see “Enabling security” on page 37.

Configuring user ID for proper privileges

Use this page to configure a user ID for proper privileges or to log on as a service on the Windows platform.

Windows

1. Click **Start > Settings > Control Panel > Administrative Tools > Local Security Policy > Local Policies > User Rights Assignments > Act as part of the operating system** (or **Log on as a service**). For a Windows domain controller, replace **Local Security Policy** with **Domain Security Policy** in the first step.

Note: If the machine is a standalone machine and not a member of a domain, you must add a `machineName\userID`, where the `userID` is the owner of the process, such as WebSphere Application Server. If you run WebSphere Application Server as a service, you can log on with `localsystem` as the service.

2. If the machine is a member of a domain, add `domainName\userID`, where the `userID` is the owner of process (such as WebSphere Application Server). Start WebSphere Application Server as a service with login ID `domainName\userID`. If WebSphere Application Server is already in service, go to the service and right-click **IBM WebSphere Application Server > properties > Log on** to change the logon ID and password.
3. Add the user name by clicking **Add**.
4. Restart the machine.

What to do next

Note: In all of the previous configurations, the server can be run as a service using LocalSystem for the Log On As entry. The LocalSystem entry has the required privileges and there is no need to give special privileges to any user. However, because the LocalSystem entry has special privileges, make sure that it is appropriate to use in your environment.

Local operating system settings

Use this page to configure local operating system registry settings.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Local operating system**.
3. Click **Configure**.

WebSphere Application Server Version 6.1 distinguishes between the user identities for administrators who manage the environment and server identities for authenticating server to server communications. In most cases, server identities are automatically generated and are not stored in a repository.

However, if you are adding a Version 5.0.x or 6.0.x node to a Version 6.1 cell, you must ensure that the Version 5.x or Version 6.0.x server identity and password are defined in the repository for this cell. Enter the server user identity and password on this panel.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your local operating system.

The user name is used to log on to the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Automatically generated server identity:

Enables the application server to generate the server identity, which is recommended for environments that contain only Version 6.1 or later nodes. Automatically generated server identities are not stored in a user repository.

You can change this server identity on the Authentication mechanisms and expiration panel. To access the Authentication mechanisms and expiration panel, click **Security > Global security > Authentication mechanisms and expiration**. Change the value of the Internal server ID field.

Default: Enabled

Server identity that is stored in the repository:

Specifies a user identity in the repository that is used for internal process communication. Cells that contain Version 5.x or 6.0.x nodes require a server user identity that is defined in the active user repository.

Default: Enabled

Server user ID or administrative user on a Version 6.0.x node:

Specifies the user ID that is used to run the application server for security purposes.

Password:

Specifies the password that corresponds to the server ID.

Local operating system wizard settings

Use this security wizard page to configure local operating system registry settings.

To view this security wizard page, complete the following steps:

1. Click **Security > Global security > Security configuration wizard**.
2. Select your protection settings and click **Next**.
3. Select the **Local operating system** option and click **Next**.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your local operating system.

The user name is used to log on to the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Configuring Lightweight Directory Access Protocol user registries

To access a user registry using the Lightweight Directory Access Protocol (LDAP), you must know a valid user name (ID) and password, the server host and port of the registry server, the base distinguished name (DN) and, if necessary, the bind DN and the bind password. You can choose any valid user in the user registry that is searchable. You can use any user ID that has the administrative role to log in.

Before you begin

In some LDAP servers, administrative users cannot be searched and thus cannot be used. The user is referred to as a WebSphere Application Server security server ID, server ID, or server user ID in the documentation. A server ID user has special privileges when calling some protected internal methods.

Normally, the primary administrative user name is used to log into the administrative console if security is enabled. By default, security is enabled after installation.

When security is enabled in the product, the primary administrative user name and password are authenticated with the registry during the product startup. If authentication fails, the server does not start. It is important to choose an ID and password that do not expire or change often. If the product server user ID or password need to change in the registry, make sure that the changes are performed when all the product servers are up and running. When changes are to be made in the registry, review the article on “Standalone Lightweight Directory Access Protocol registries” on page 242 (LDAP) before beginning this task.

1. In the administrative console, click **Security > Global security**.

2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
3. Enter a valid user name in the **Primary administrative user name** field. You can either enter the complete distinguished name (DN) of the user or the short name of the user, as defined by the user filter in the Advanced LDAP settings panel. For example, enter the user ID for Netscape browsers. This ID is the security server ID, which is only used for WebSphere Application Server security and is not associated with the system process that runs the server. The server calls the local operating system registry to authenticate and obtain privilege information about users by calling the native application programming interfaces (API) in that particular registry.
4. Optional: If you want to use the server ID that is stored in the repository, complete the following:
 - a. Select **Automatically generated server identity** to enable the application server to generate the server identity that is used for internal process communication. You can change this server identity on the Authentication mechanisms and expiration panel. To access the Authentication mechanisms and expiration panel, click **Security > Global security > Authentication mechanisms and expiration**. Change the value of the **Internal server ID** field.
 - b. Alternatively, specify a user identity in the repository that is used for internal process communication in the **Server identity that is stored in the repository** field.
 - c. Alternatively, specify the user ID that is used to run the application server for security purposes in the **Server user ID or administrative user on a Version 6.0.x node** field.
5. Select the type of LDAP server to use from the **Type** list. The type of LDAP server determines the default filters that are used by WebSphere Application Server. These default filters change the **Type** field to **Custom**, which indicates that custom filters are used. This action occurs after you click **OK** or **Apply** in the Advanced LDAP settings panel. Choose the **Custom** type from the list and modify the user and group filters to use other LDAP servers, if required.

IBM Tivoli Directory Server users can choose IBM Tivoli Directory Server as the directory type. Use the IBM Tivoli Directory Server directory type for better performance. For a list of supported LDAP servers, see the Supported hardware, software, and APIs Web site.

Note: IBM SecureWay® Directory Server has been renamed to IBM Tivoli Directory Server in WebSphere Application Server version 6.1.

6. Enter the fully qualified host name of the LDAP server in the **Host** field. You can enter either the IP address or domain name system (DNS) name.
7. Enter the LDAP server port number in the **Port** field. The host name and the port number represent the realm for this LDAP server in the WebSphere Application Server cell. So, if servers in different cells are communicating with each other using Lightweight Third Party Authentication (LTPA) tokens, these realms must match exactly in all the cells.

The default value is 389. If multiple WebSphere Application Servers are installed and configured to run in the same single sign-on domain, or if the WebSphere Application Server interoperates with a previous version of the WebSphere Application Server, then it is important that the port number match all configurations. For example, if the LDAP port is explicitly specified as 389 in a version 5.x configuration, and a WebSphere Application Server at version 6.0.x is going to interoperate with the version 5.x server, then verify that port 389 is specified explicitly for the version 6.0.x server.
8. Enter the base distinguished name (DN) in the **Base distinguished name** field. The base DN indicates the starting point for searches in this LDAP directory server. For example, for a user with a DN of cn=John Doe, ou=Rochester, o=IBM, c=US, specify the base DN as any of the following options assuming a suffix of c=us): ou=Rochester, o=IBM, c=us or o=IBM c=us or c=us. For authorization purposes, this field is case sensitive by default. Match the case in your directory server. If a token is received (for example, from another cell or Lotus Domino) the base DN in the server must match exactly the base DN from the other cell or Domino. If case sensitivity is not a consideration for authorization, enable the **Ignore case for authorization** option.

In WebSphere Application Server, the distinguished name is normalized according to the Lightweight Directory Access Protocol (LDAP) specification. Normalization consists of removing spaces in the base distinguished name before or after commas and equal symbols. An example of a

non-normalized base distinguished name is `o = ibm, c = us` or `o=ibm, c=us`. An example of a normalized base distinguished name is `o=ibm,c=us`.

To interoperate between WebSphere Application Server Version 5 and later versions, you must enter a normalized base distinguished name in the **Base Distinguished Name** field. In WebSphere Application Server, Version 5.0.1 or later, the normalization occurs automatically during runtime.

This field is required for all LDAP directories except the Lotus Domino Directory. The **Base Distinguished Name** field is optional for the Domino server.

9. Optional: Enter the bind DN name in the **Bind distinguished name** field. The bind DN is required if anonymous binds are not possible on the LDAP server to obtain user and group information. If the LDAP server is set up to use anonymous binds, leave this field blank. If a name is not specified, the application server binds anonymously. See the **Base Distinguished Name** field description for examples of distinguished names.
10. Optional: Enter the password corresponding to the bind DN in the **Bind password** field.
11. Optional: Modify the Search time out value. This timeout value is the maximum amount of time that the LDAP server waits to send a response to the product client before stopping the request. The default is 120 seconds.
12. Ensure that the **Reuse connection** option is selected. This option specifies that the server should reuse the LDAP connection. Clear this option only in rare situations where a router is used to send requests to multiple LDAP servers and when the router does not support affinity. Leave this option selected for all other situations.
13. Optional: Verify that the **Ignore case for authorization** option is enabled. When you enable this option, the authorization check is case insensitive. Normally, an authorization check involves checking the complete DN of a user, which is unique in the LDAP server and is case sensitive. However, when you use either the IBM Directory Server or the Sun ONE (formerly iPlanet) Directory Server LDAP servers, you must enable this option because the group information that is obtained from the LDAP servers is not consistent in case. This inconsistency affects the authorization check only. Otherwise, this field is optional and can be enabled when a case sensitive authorization check is required. For example, you might select this option when you use certificates and the certificate contents do not match the case of the entry in the LDAP server.

You can also enable the **Ignore case for authorization** option when you are using single sign-on (SSO) between the product and Lotus Domino. The default is enabled.

14. Optional: Select the **SSL enabled** option if you want to use Secure Sockets Layer communications with the LDAP server.

If you select the **SSL enabled** option, you can select either the **Centrally managed** or the **Use specific SSL alias** option.

Centrally managed

Enables you to specify an SSL configuration for particular scope such as the cell, node, server, or cluster in one location. To use the **Centrally managed** option, you must specify the SSL configuration for the particular set of endpoints. The Manage endpoint security configurations and trust zones panel displays all of the inbound and outbound endpoints that use the SSL protocol. If you expand the Inbound or Outbound section of the panel and click the name of a node, you can specify an SSL configuration that is used for every endpoint on that node. For an LDAP registry, you can override the inherited SSL configuration by specifying an SSL configuration for LDAP. To specify an SSL configuration for LDAP, complete the following steps:

- a. Click **Security > SSL certificate and key management > Manage endpoint security configurations and trust zones**.
- b. Expand **Outbound > cell_name > Nodes > node_name > Servers > server_name > LDAP**.

Use specific SSL alias

Select the **Use specific SSL alias** option if you intend to select one of the SSL configurations in the menu below the option.

This configuration is used only when SSL is enabled for LDAP. The default is *DefaultSSLSettings*. To modify or create a new SSL configuration, complete the following steps:

- a. Click **Security > SSL certificate and key management**.
 - b. Under Configuration settings, click **Manage endpoint security configurations**.
 - c. Select a Secure Sockets Layer (SSL) *configuration_name* for selected scopes, such as a cell, node, server, or cluster.
 - d. Under Related items, click **SSL configurations**.
 - e. Click **New**.
15. Click **OK** and either **Apply** or **Save** until you return to the Global security panel.

Results

This set of steps is required to set up the LDAP user registry. This step is required as part of enabling security in the WebSphere Application Server.

What to do next

1. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53.
2. Save, stop, and restart all the product servers (deployment managers, nodes and Application Servers) for changes in this panel to take effect. If the server comes up without any problems the setup is correct.

Standalone LDAP registry settings

Use this page to configure Lightweight Directory Access Protocol (LDAP) settings when users and groups reside in an external LDAP directory.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.

When security is enabled and any of these properties change, go to the Global security panel and click **Apply** to validate the changes.

WebSphere Application Server Version 6.1 distinguishes between the user identities for administrators who manage the environment and server identities for authenticating server to server communications. In most cases, server identities are automatically generated and are not stored in a repository.

However, if you are adding a Version 5.0.x or 6.0.x node to a Version 6.1 cell, you must ensure that the Version 5.x or Version 6.0.x server identity and password are defined in the repository for this cell. Enter the server user identity and password on this panel.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your custom user registry.

The user name is used to log onto the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Automatically generated server identity:

Enables the application server to generate the server identity, which is recommended for environments that contain only Version 6.1 or later nodes. Automatically generated server identities are not stored in a user repository.

You can change this server identity on the Authentication mechanisms and expiration panel. To access the Authentication mechanisms and expiration panel, click **Security > Global security > Authentication mechanisms and expiration**. Change the value of the Internal server ID field.

Default: Enabled

Server identity that is stored in the repository:

Specifies a user identity in the repository that is used for internal process communication. Cells that contain Version 5.x or 6.0.x nodes require a server user identity that is defined in the active user repository.

Default: Enabled

Server user ID or administrative user on a Version 6.0.x node:

Specifies the user ID that is used to run the application server for security purposes.

Password:

Specifies the password that corresponds to the server ID.

Type of LDAP server:

Specifies the type of LDAP server to which you connect.

IBM SecureWay Directory Server is not supported.

Host:

Specifies the host ID (IP address or domain name service (DNS) name) of the LDAP server.

Port:

Specifies the host port of the LDAP server.

If multiple application servers are installed and configured to run in the same single sign-on domain or if the application server interoperates with a previous version, it is important that the port number match all configurations. For example, if the LDAP port is explicitly specified as 389 in a Version 4.0.x configuration, and a WebSphere Application Server at Version 5 is going to interoperate with the Version 4.0.x server, verify that port 389 is specified explicitly for the Version 5 server.

Default: 389
Type: Integer

Base distinguished name (DN):

Specifies the base distinguished name (DN) of the directory service, which indicates the starting point for LDAP searches of the directory service. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed.

For example, for a user with a DN of cn=John Doe , ou=Rochester, o=IBM, c=US, specify the Base DN as any of the following options: ou=Rochester, o=IBM, c=US or o=IBM c=US or c=US. For authorization purposes, this field is case sensitive. This specification implies that if a token is received, for example, from another cell or Lotus Domino, the base DN in the server must match the base DN from the other cell or Lotus Domino server exactly. If case sensitivity is not a consideration for authorization, enable the **Ignore case for authorization** option. This option is required for all Lightweight Directory Access Protocol (LDAP) directories, except for the Lotus Domino Directory, IBM Tivoli Directory Server V6.0, and Novell eDirectory, where this field is optional.

If you need to interoperate between the application server Version 5 and a Version 5.0.1 or later server, you must enter a normalized base DN. A normalized base DN does not contain spaces before or after commas and equal symbols. An example of a non-normalized base DN is o = ibm, c = us or o=ibm, c=us. An example of a normalized base DN is o=ibm,c=us. In WebSphere Application Server, Version 5.0.1 or later, the normalization occurs automatically during runtime.

Bind distinguished name (DN):

Specifies the DN for the application server to use when binding to the directory service.

If no name is specified, the application server binds anonymously. See the Base distinguished name (DN) field description for examples of distinguished names.

Bind password:

Specifies the password for the application server to use when binding to the directory service.

Search timeout:

Specifies the timeout value in seconds for a Lightweight Directory Access Protocol (LDAP) server to respond before stopping a request.

Default: 120

Reuse connection:

Specifies whether the server reuses the LDAP connection. Clear this option only in rare situations where a router is used to distribute requests to multiple LDAP servers and when the router does not support affinity.

Default: Enabled
Range: Enabled or Disabled

Note: Disabling the **Reuse connection** option causes the application server to create a new LDAP connection for every LDAP search request. This situation impacts system performance if your environment requires extensive LDAP calls. This option is provided because the router is not sending the request to the same LDAP server. The option is also used when the idle connection timeout value or firewall timeout value between the application server and LDAP is too small.

If you are using WebSphere Edge Server for LDAP failover, you must enable TCP resets with the Edge server. A TCP reset causes the connection to immediately closed and a backup server to failover. For more information, see "Sending TCP resets when server is down" at <http://www.ibm.com/software/webservers/appserv/doc/v50/ec/infocenter/edge/LBguide.htm#HDRRESETSERVER> and the Edge Server V2 - TCP Reset feature in PTF #2 described in: <ftp://ftp.software.ibm.com/software/websphere/edgeserver/info/doc/v20/en/updates.pdf>.

Ignore case for authorization:

Specifies that a case insensitive authorization check is performed when using the default authorization.

This option is required when IBM Tivoli Directory Server is selected as the LDAP directory server.

This option is required when Sun ONE Directory Server is selected as the LDAP directory server. For more information, see "Using specific directory servers as the LDAP server" in the documentation.

This option is optional and can be enabled when a case-sensitive authorization check is required. For example, use this option when the certificates and the certificate contents do not match the case that is used for the entry in the LDAP server. You can enable the **Ignore case for authorization** option when using single sign-on (SSO) between the application server and Lotus Domino.

Default: Enabled
Range: Enabled or Disabled

SSL enabled:

Specifies whether secure socket communication is enabled to the Lightweight Directory Access Protocol (LDAP) server.

When enabled, the LDAP Secure Sockets Layer (SSL) settings are used, if specified.

Centrally managed:

Specifies that the selection of an SSL configuration is based upon the outbound topology view for the Java Naming and Directory Interface (JNDI) platform.

Centrally managed configurations support one location to maintain SSL configurations rather than spreading them across the configuration documents.

Default: Enabled

Use specific SSL alias:

Specifies the SSL configuration alias to use for LDAP outbound SSL communications.

This option overrides the centrally managed configuration for the JNDI platform.

Standalone LDAP registry wizard settings

Use this security wizard page to provide the basic settings to connect the application server to an existing Lightweight Directory Access Protocol (LDAP) registry.

To view this security wizard page, click **Security > Global security > Security configuration wizard**. You can modify your LDAP registry configuration by completing the following steps:

1. Click **Security > Global security**.

2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your custom user registry.

The user name is used to log onto the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Type of LDAP server:

Specifies the type of LDAP server to which you connect.

IBM SecureWay Directory Server is not supported.

Host:

Specifies the host ID (IP address or domain name service (DNS) name) of the LDAP server.

Port:

Specifies the host port of the LDAP server.

If multiple application servers are installed and configured to run in the same single sign-on domain or if the application server interoperates with a previous version, it is important that the port number match all configurations. For example, if the LDAP port is explicitly specified as 389 in a Version 4.0.x configuration, and a WebSphere Application Server at Version 5 is going to interoperate with the Version 4.0.x server, verify that port 389 is specified explicitly for the Version 5 server.

Default:	389
Type:	Integer

Base distinguished name (DN):

Specifies the base distinguished name (DN) of the directory service, which indicates the starting point for LDAP searches of the directory service. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed.

For example, for a user with a DN of `cn=John Doe , ou=Rochester, o=IBM, c=US`, specify the Base DN as any of the following options: `ou=Rochester, o=IBM, c=US` or `o=IBM, c=US` or `c=US`. For authorization purposes, this field is case sensitive. This specification implies that if a token is received, for example, from another cell or Lotus Domino, the base DN in the server must match the base DN from the other cell or Lotus Domino server exactly.

If you need to interoperate between the application server Version 5 and a Version 5.0.1 or later server, you must enter a normalized base DN. A normalized base DN does not contain spaces before or after commas and equal symbols. An example of a non-normalized base DN is `o = ibm, c = us` or `o=ibm,`

c=us. An example of a normalized base DN is o=ibm,c=us. In WebSphere Application Server, Version 5.0.1 or later, the normalization occurs automatically during run time.

Bind distinguished name (DN):

Specifies the DN for the application server to use when binding to the directory service.

If no name is specified, the application server binds anonymously. See the Base distinguished name (DN) field description for examples of distinguished names.

Bind password:

Specifies the password for the application server to use when binding to the directory service.

Advanced Lightweight Directory Access Protocol user registry settings

Use this page to configure the advanced Lightweight Directory Access Protocol (LDAP) user registry settings when users and groups reside in an external LDAP directory.

To view this administrative page, complete the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
3. Under Additional properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.

Default values for all the user and group related filters are already completed in the appropriate fields. You can change these values depending on your requirements. These default values are based on the type of LDAP server that is selected in the Standalone LDAP registry settings panel. If this type changes, for example from Netscape to Secureway, the default filters automatically change. When the default filter values change, the LDAP server type changes to Custom to indicate that custom filters are used. When security is enabled and any of these properties change, go to the Global security panel and click **Apply** or **OK** to validate the changes.

User filter:

Specifies the LDAP user filter that searches the user registry for users.

This option is typically used for security role-to-user assignments and specifies the property by which to look up users in the directory service. For example, to look up users based on their user IDs, specify (&(uid=%v)(objectclass=inetOrgPerson)). For more information about this syntax, see the LDAP directory service documentation.

Data type: String

Group filter:

Specifies the LDAP group filter that searches the user registry for groups

This option is typically used for security role-to-group assignments and specifies the property by which to look up groups in the directory service. For more information about this syntax, see the LDAP directory service documentation.

Data type: String

User ID map:

Specifies the LDAP filter that maps the short name of a user to an LDAP entry.

Specifies the piece of information that represents users when users display. For example, to display entries of the object class = inetOrgPerson type by their IDs, specify inetOrgPerson:uid. This field takes multiple objectclass:property pairs delimited by a semicolon (;).

Data type: String

Group ID map:

Specifies the LDAP filter that maps the short name of a group to an LDAP entry.

Specifies the piece of information that represents groups when groups display. For example, to display groups by their names, specify *:cn. The asterisk (*) is a wildcard character that searches on any object class in this case. This field takes multiple objectclass:property pairs, delimited by a semicolon (;).

Data type: String

Group member ID map:

Specifies the LDAP filter that identifies user-to-group relationships.

For directory types SecureWay, and Domino, this field takes multiple objectclass:property pairs, delimited by a semicolon (;). In an objectclass:property pair, the object class value is the same object class that is defined in the group filter, and the property is the member attribute. If the object class value does not match the object class in the group filter, authorization might fail if groups are mapped to security roles. For more information about this syntax, see your LDAP directory service documentation.

For IBM Directory Server, Sun ONE, and Active Directory, this field takes multiple group attribute:member attribute pairs delimited by a semicolon (;). These pairs are used to find the group memberships of a user by enumerating all the group attributes that are possessed by a given user. For example, attribute pair memberof:member is used by Active Directory, and ibm-allGroup:member is used by IBM Directory Server. This field also specifies which property of an object class stores the list of members belonging to the group represented by the object class. For supported LDAP directory servers, see "Supported directory services".

Data type: String

Perform a nested group search:

Specifies a recursive nested group search.

Select this option if the Lightweight Directory Access Protocol (LDAP) server does not support recursive server-side group member searches and if recursive group member search is required. It is not recommended that you select this option to locate recursive group memberships for LDAP servers. Application server security leverages the recursive search functionality of the LDAP server to search a user's group memberships, including recursive group memberships. For example:

- IBM Directory Server is preconfigured by the application server security to recursively calculate a user's group memberships using the ibm-allGroup attribute.
- SunONE directory server is preconfigured to calculate nested group memberships using the nsRole attribute.

Data type: String

Kerberos user filter:

Specifies the Kerberos user filter value. This value can be modified when Kerberos is configured and is active as one of the preferred authentication mechanisms.

Data type: String

Certificate map mode:

Specifies whether to map X.509 certificates into an LDAP directory by EXACT_DN or CERTIFICATE_FILTER. Specify CERTIFICATE_FILTER to use the specified certificate filter for the mapping.

Data type: String

Certificate filter:

Specifies the filter certificate mapping property for the LDAP filter. The filter is used to map attributes in the client certificate to entries in the LDAP registry.

If more than one LDAP entry matches the filter specification at runtime, authentication fails because the result is an ambiguous match. The syntax or structure of this filter is:

(&(uid=\${SubjectCN})(objectclass=inetOrgPerson)). The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. The right side must begin with a dollar sign (\$) and open bracket ({} and end with a close bracket (}). You can use the following certificate attribute values on the right side of the filter specification. The case of the strings is important:

- \${UniqueKey}
- \${PublicKey}
- \${IssuerDN}
- \${Issuer<xx>}

where <xx> is replaced by the characters that represent any valid component of the Issuer Distinguished Name. For example, you might use \${IssuerCN} for the Issuer Common Name.

- \${NotAfter}
- \${NotBefore}
- \${SerialNumber}
- \${SigAlgName}
- \${SigAlgOID}
- \${SigAlgParams}
- \${SubjectDN}
- \${Subject<xx>}

where <xx> is replaced by the characters that represent any valid component of the Subject Distinguished Name. For example, you might use \${SubjectCN} for the Subject Common Name.

- \${Version}

Data type: String

Configuring Lightweight Directory Access Protocol search filters

Use this topic to configure the LDAP search filters. These steps are required to modify existing user and group filters for a particular LDAP directory type, and also to set up certificate filters to map certificates to entries in the LDAP server.

Before you begin

WebSphere Application Server uses Lightweight Directory Access Protocol (LDAP) filters to search and obtain information about users and groups from an LDAP directory server. A default set of filters is provided for each LDAP server that the product supports. You can modify these filters to fit your LDAP configuration. After the filters are modified and you click **OK** or **Apply** the directory type in the Standalone LDAP registry panel changes to *custom*, which indicates that custom filters are used. Also, you can develop filters to support any additional type of LDAP server. The effort to support additional LDAP directories is optional and other LDAP directory types are not supported. Complete the following steps in the administrative console.

1. Click **Security > Global security**.
2. Under User account repository, select **Standalone LDAP registry** and click **Configure**.
3. Under Additional properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
4. Modify the user filter, if necessary. The user filter is used for searching the registry for users and is typically used for the security role-to-user assignment. The filter is also used to authenticate a user with the attribute that is specified in the filter. The filter specifies the property that is used to look up users in the directory service.

In the following example, the property that is assigned to %v, which is the short name of the user, must be a unique key. Two LDAP entries with the same object class cannot have the same short name. To look up users based on their user IDs (uid) and to use the inetOrgPerson object class, specify the following syntax:

```
(&(uid=%v)(objectclass=inetOrgPerson))
```

For more information about this syntax, see the “Using specific directory servers as the LDAP server” on page 133 documentation.

5. Modify the Kerberos user filter, if necessary. The Kerberos user filter name is used for searching the registry for the Kerberos principal name. Specify the LDAP attribute that holds the Kerberos principal name.

IBM Lotus Domino default krbuser filter:

```
(&(krbPrincipalName=%v)(objectcategory=Person))
```

IBM SecureWay Directory Server default krbuser filter:

```
(&(krbPrincipalName=%v)(objectcategory=ePerson))
```

Microsoft Active Directory default krbuser filter:

```
(&(userprincipalname=%v)(objectcategory=user))
```

Sun Java System Directory Server default krbuser filter:

```
(&(krbPrincipalName=%v)(objectcategory=inetOrgPerson))
```

Novell eDirectory default krbuser filter:

```
(&(krbPrincipalName=%v)(objectcategory=Person))
```

6. Optional: If your using Federated Repositories, modify the Kerberos attribute name if necessary. The Kerberos attribute name is used for searching the registry for Kerberos principal. Specify the LDAP attribute that holds the Kerberos principal name.

IBM Lotus Domino default krbuser filter:

```
krbPrincipalName
```

IBM SecureWay Directory Server default krbuser filter:

```
krbPrincipalName
```

Microsoft Active Directory default krbuser filter:

```
userprincipalname
```


Sun Java System Directory Server default krbuser filter:

krbPrincipalName

Novell eDirectory default krbuser filter:

krbPrincipalName

7. Modify the group filter, if necessary. The group filter is used in searching the registry for groups and is typically used for the security role-to-group assignment. Also, the filter is used to specify the property by which to look up groups in the directory service.

In the following example, the property that is assigned to %v, which is the short name of the group, must be a unique key. Two LDAP entries with the same object class cannot have the same short name. To look up groups based on their common names (CN) and to use either the groupOfNames object class or the groupOfUniqueNames object class, specify the following syntax:

```
(&(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames)))
```

For more information about this syntax, see the “Using specific directory servers as the LDAP server” on page 133 documentation.

8. Modify the user ID map, if necessary. This filter maps the short name of a user to an LDAP entry and specifies the piece of information that represents users when these users are displayed with their short names. For example, to display entries of object class = inetOrgPerson by their IDs, specify inetOrgPerson:uid. This field takes multiple objectclass:property pairs, delimited by a semicolon (;). To provide a consistent value for methods like the getCallerPrincipal method and the getUserPrincipal method, the short name that is obtained by using this filter is used. For example, the CN=Bob Smith, ou=austin.ibm.com, o=IBM, c=US user can log in using any attributes that are defined, for example, e-mail address, social security number, and so on, but when these methods are called, the bob user ID is returned no matter how the user logs in.
9. Modify the group ID map filter, if necessary. This filter maps the short name of a group to an LDAP entry and specifies the piece of information that represents groups when groups display. For example, to display groups by their names, specify *:cn. The asterisk (*) is a wildcard character that searches on any object class in this case. This field takes multiple objectclass:property pairs, delimited by a semicolon (;).
10. Modify the group member ID map filter, if necessary. This filter identifies user-to-group memberships. For SecureWay, and Domino directory types, this field is used to query all the groups that match the specified object classes to see if the user is contained in the specified attribute. For example, to get all the users that belong to groups with the groupOfNames object class and the users that are contained in the member attributes, specify groupOfNames:member. This syntax, which is a property of an object class, stores the list of members that belong to the group that is represented by the object class. This field takes multiple objectclass:property pairs that are delimited by a semicolon (;). For more information about this syntax, see the “Using specific directory servers as the LDAP server” on page 133.
For the IBM Tivoli Directory Server, Sun ONE, and Active Directory, this field is used to query all users in a group with the information that is stored in the user object. For example, the memberof:member filter (for Active Directory) is used to get the memberof attribute of the user object to obtain all the groups to which the user belongs. The member attribute is used to get all the users in a group that use the Group object. Using the User object to obtain the group information improves performance.
11. Select the **Perform a nested group search** option if your LDAP server does not support recursive server-side searches.
12. Modify the Certificate map mode, if necessary. You can use the X.590 certificates for user authentication when LDAP is selected as the registry. This field is used to indicate whether to map the X.509 certificates into an LDAP directory user by **EXACT_DN** or **CERTIFICATE_FILTER**. If **EXACT_DN** is selected, the DN in the certificate must exactly match the user entry in the LDAP server, including case and spaces.

Select the **Ignore case for authorization** option on the Standalone LDAP registry settings to make the authorization case insensitive. To access the Standalone LDAP registry settings panel, complete the following steps:

- a. Click **Security > Global security**.
 - b. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**.
13. If you select **CERTIFICATE_FILTER**, specify the LDAP filter for mapping attributes in the client certificate to entries in LDAP. If more than one LDAP entry matches the filter specification at run time, authentication fails because an ambiguous match results. The syntax or structure of this filter is: LDAP attribute=\${Client certificate attribute} (for example, uid=\${SubjectCN}).

The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. Note that the right side must begin with a dollar sign (\$), open bracket ({}), and end with a close bracket (}). Use the following certificate attribute values on the right side of the filter specification. The case of the strings is important.

- \${UniqueKey}
- \${PublicKey}
- \${IssuerDN}
- \${Issuer<xx>}

where <xx> is replaced by the characters that represent any valid component of the Issuer Distinguished Name. For example, you might use \${IssuerCN} for the Issuer Common Name.

- \${NotAfter}
- \${NotBefore}
- \${SerialNumber}
- \${SigAlgName}
- \${SigAlgOID}
- \${SigAlgParams}
- \${SubjectDN}
- \${Subject<xx>}

where <xx> is replaced by the characters that represent any valid component of the Subject Distinguished Name. For example, you might use \${SubjectCN} for the Subject Common Name.

- \${Version}

To enable this field, select **CERTIFICATE_FILTER** for the certificate mapping.

14. Click **Apply**.

When any LDAP user or group filter is modified in the Advanced LDAP Settings panel click **Apply**. Clicking **OK** navigates you to the Standalone LDAP registry panel, which contains the previous LDAP directory type, rather than the custom LDAP directory type. Clicking **OK** or **Apply** in the Standalone LDAP registry panel saves the back-level LDAP directory type and the default filters of that directory. This action overwrites any changes to the filters that you made. To avoid overwriting changes, you can take either of the following actions:

- Click **Apply** in the Advanced Lightweight Directory Access Protocol (LDAP) user registry settings panel. Click **Security > Global security** and change the User account repository type to Standalone custom registry.
- Select **Custom** type from the Standalone LDAP registry panel. Click **Apply** and then change the filters by clicking the Advanced Lightweight Directory Access Protocol (LDAP) user registry settings panel. After you complete your changes, click **Apply** or **OK**.

The validation of the changes does not take place in this panel. Validation is done when you click **OK** or **Apply** on the Global security panel. If you are in the process of enabling security for the first time, complete the remaining steps and go to the Global security panel. Select **Standalone LDAP registry** as the user account repository. If security is already enabled and any information on this panel changes, go to the Global security panel and click **OK** or **Apply** to validate your changes. If your changes are not validated, the server might not start.

Results

These steps result in the configuration of the LDAP search filters. These steps are required to modify existing user and group filters for a particular LDAP directory type. The steps are also used to set up certificate filters to map certificates to entries in the LDAP server.

What to do next

1. Validate this setup by clicking **OK** or **Apply** on the Global security panel.
2. Save, stop, and start all the product servers, including the cell, nodes and all of the application servers for any changes in this panel to become effective.
3. After the server starts, go through all the security-related tasks (getting users, getting groups, and so on) to verify that the changes to the filters function.

Using specific directory servers as the LDAP server

This article provides important information about the directory servers that are supported as Lightweight Directory Access Protocol (LDAP) servers in WebSphere Application Server.

Before you begin

Microsoft® Active Directory forest is not supported in the user registry in this product.

About this task

For a list of supported LDAP servers, refer to the Supported hardware and software Web site.

It is expected that other LDAP servers follow the LDAP specification. Support is limited to these specific directory servers only. You can use any other directory server by using the custom directory type in the list and by filling in the filters that are required for that directory.

To improve performance for LDAP searches, the default filters for IBM Tivoli Directory Server, Sun ONE, and Active Directory are defined such that when you search for a user, the result contains all the relevant information about the user (user ID, groups, and so on). As a result, the product does not call the LDAP server multiple times. This definition is possible only in these directory types, which support searches where the complete user information is obtained.

If you use the IBM Directory Server, select the **Ignore case for authorization** option. This option is required because when the group information is obtained from the user object attributes, the case is not the same as when you get the group information directly. For the authorization to work in this case, perform a case insensitive check and verify the requirement for the **Ignore case for authorization** option.

- **Using IBM Tivoli Directory Server as the LDAP server**

To use IBM Tivoli Directory Server, formerly IBM Directory Server, select **IBM Tivoli Directory Server** as the directory type.

The difference between these two types is group membership lookup. It is recommended that you choose the IBM Tivoli Directory Server for optimum performance during runtime. In the IBM Tivoli Directory Server, the group membership is an operational attribute. With this attribute, a group membership lookup is done by enumerating the `ibm-allGroups` attribute for the entry. All group memberships, including the static groups, dynamic groups, and nested groups, can be returned with the `ibm-allGroups` attribute.

WebSphere Application Server supports dynamic groups, nested groups, and static groups in IBM Tivoli Directory Server using the `ibm-allGroups` attribute. To utilize this attribute in a security authorization application, use a case-insensitive match so that attribute values returned by the `ibm-allGroups` attribute are all in uppercase.

Note: It is recommended that you do not install IBM Tivoli Directory Server Version 6.0 on the same machine that you install WebSphere Application Server Version 7.0. IBM Tivoli Directory Server

Version 6.0 includes WebSphere Application Server Express Version 5.1.1, which the directory server uses for its administrative console. Install the Web Administration tool Version 6.0 and WebSphere Application Server Express Version 5.1.1, which are both bundled with IBM Tivoli Directory Server Version 6.0, on a different machine from WebSphere Application Server Version 7.0. You cannot use WebSphere Application Server Version 7.0 as the administrative console for IBM Tivoli Directory Server. If IBM Tivoli Directory Server Version 6.0 and WebSphere Application Server Version 7.0 are installed on the same machine, you might encounter port conflicts.

If you must install IBM Tivoli Directory Server Version 6.0 and WebSphere Application Server Version 7.0 on the same machine, consider the following information:

- During the IBM Tivoli Directory Server installation process, you must select both the **Web Administration tool** and **WebSphere Application Server Express Version 5.1.1**.
- Install WebSphere Application Server Version 7.0.
- When you install WebSphere Application Server Version 7.0, change the port number for the application server.
- You might need to adjust the WebSphere Application Server environment variables on WebSphere Application Server Version 7.0 for *WAS_HOME* and *WAS_INSTALL_ROOT* (or *APP_SERVER_ROOT* for i5/OS®). To change the variables using the administrative console, click **Environment > WebSphere Variables**.

- **Using a Lotus Domino Enterprise Server as the LDAP server**

If you select the Lotus Domino Enterprise Server Version 6.5.4 or Version 7.0 and the attribute short name is not defined in the schema, you can take either of the following actions:

- Change the schema to add the short name attribute.
- Change the user ID map filter to replace the short name with any other defined attribute (preferably to UID). For example, change `person:shortname` to `person:uid`.

The userID map filter is changed to use the uid attribute instead of the shortname attribute as the current version of Lotus Domino does not create the shortname attribute by default. If you want to use the shortname attribute, define the attribute in the schema and change the userID map filter.

User ID Map : person:shortname

- **Using Sun ONE Directory Server as the LDAP server**

You can select **Sun ONE Directory Server** for your Sun ONE Directory Server system. In Sun ONE Directory Server, the object class is the default `groupOfUniqueName` when you create a group. For better performance, WebSphere Application Server uses the User object to locate the user group membership from the `nsRole` attribute. Create the group from the role. If you want to use the `groupOfUniqueName` attribute to search groups, specify your own filter setting. Roles unify entries. Roles are designed to be more efficient and easier to use for applications. For example, an application can locate the role of an entry by enumerating all the roles that are possessed by a given entry, rather than selecting a group and browsing through the members list. When using roles, you can create a group using a:

- Managed role
- Filtered role
- Nested role

All of these roles are computable by the `nsRole` attribute.

- **Using Microsoft Active Directory server as the LDAP server**

To use Microsoft Active Directory as the LDAP server for authentication with WebSphere Application Server you must take specific steps. By default, Microsoft Active Directory does not permit anonymous LDAP queries. To create LDAP queries or to browse the directory, an LDAP client must bind to the LDAP server using the distinguished name (DN) of an account that has the authority to search and read the values of LDAP attributes, such as user and group information, needed by the Application Server. A group membership search in the Active Directory is done by enumerating the `memberof` attribute for a

given user entry, rather than browsing through the member list in each group. If you change the default behavior to browse each group, you can change the **Group Member ID Map** field from memberof:member to group:member.

The following steps describe how to set up Microsoft Active Directory as your LDAP server.

1. Determine the full distinguished name (DN) and password of an account in the administrators group. For example, if the Active Directory administrator creates an account in the Users folder of the Active Directory Users and Computers Windows control panel and the DNS domain is ibm.com, the resulting DN has the following structure:

```
cn=<adminUsername>, cn=users, dc=ibm,  
dc=com
```
2. Determine the short name and password of any account in the Microsoft Active Directory.
3. Use the WebSphere Application Server administrative console to set up the information that is needed to use Microsoft Active Directory.
 - a. Click **Security > Global security**.
 - b. Under User account repository, select **Standalone LDAP registry** and click **Configure**.
 - c. Set up LDAP with Active Directory as the type of LDAP server. Based on the information that is determined in the previous steps, you can specify the following values on the LDAP settings panel:

Primary administrative user name

Specify the name of a user with administrative privileges that is defined in the registry. This user name is used to access the administrative console or used by wsadmin.

Type Specify Active Directory

Host Specify the domain name service (DNS) name of the machine that is running Microsoft Active Directory.

Base distinguished name (DN)

Specify the domain components of the DN of the account that is chosen in the first step. For example: dc=ibm, dc=com

Bind distinguished name (DN)

Specify the full distinguished name of the account that is chosen in the first step. For example: cn=*adminUsername*, cn=users, dc=ibm, dc=com

Bind password

Specify the password of the account that is chosen in the first step.

- d. Click **OK** and **Save** to save the changes to the master configuration.
4. Click **Security > Global security**.
 5. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
 6. Select either the **Automatically generated server identity** or **Server identity that is stored in the repository** option. If you select the **Server identity that is stored in the repository** option, enter the following information:

Server user ID or administrative user on a Version 6.0.x node

Specify the short name of the account that is chosen in the second step.

Server user password

Specify the password of the account that is chosen in the second step.

7. Optional: Set ObjectCategory as the filter in the Group member ID map field to improve LDAP performance.
 - a. Under Additional properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
 - b. Add ;objectCategory:group to the end of the Group member ID map field.

8. Click **OK** and **Save** to save the changes to the master configuration.
9. Stop and restart the administrative server so that the changes take effect.

Locating a user's group memberships in Lightweight Directory Access Protocol

WebSphere Application Server security can be configured to search group memberships directly or indirectly. It can also be configured to search only a static group, or it can be configured to search static groups, recursive or nested groups, and dynamic groups for some Lightweight Directory Access Protocol (LDAP) servers.

- Evaluate group memberships from user object directly.
 - Several popular LDAP servers enable user objects to contain information about the groups to which they belong such as Microsoft Active Directory Server, or eDirectory. Some user group memberships can be computable attributes from the user object such as IBM Directory Server or Sun ONE directory server. In some LDAP servers, this attribute can be used to include a user's dynamic group memberships, nesting group memberships, and static group memberships to locate all the group memberships from a single attribute.
 - For example, in IBM Directory Server all group memberships including the static groups, dynamic groups, and nested groups can be returned using the `ibm-allGroups` attribute. In Sun ONE, all roles, including managed roles, filtered roles, and nested roles, are calculated using the `nsRole` attribute. If an LDAP server has such an attribute in a User object to include dynamic groups, nested groups, and static groups, WebSphere Application Server security can be configured to use this attribute to support these groups.
- Evaluate group memberships from a Group object indirectly.
 - Some LDAP servers enable only Group objects, such as the Lotus Domino LDAP server to contain information about users. The LDAP server does not enable the User object to contain information about groups. For this type of LDAP server, group membership searches are performed by locating the user on the member list of groups. The member list evaluation is not currently used in the static group membership search for WebSphere Application Server.
- Use the direct method for searching group memberships if your LDAP server has an attribute in the User object to include group information. To use the direct method or the indirect method, enter the appropriate value in the Group Member ID map field on the Advanced LDAP Settings panel using the following methods.
 - `objectclass:attribute` pairs for the indirect method
 - `attribute:attribute` pairs for the direct method
- Use the sample entries of `attribute:attribute` pairs in Group member ID map fields. Note that the `groupMembership` attribute lists all the static groups for which a user is a member. This attribute is NOT updated whenever an object matches or does not match a dynamic group's filter. Please refer to the Novell eDirectory documentation for more information about the `groupMembership` attribute.
 - `ibm-allGroups:member` for IBM Directory server
 - `nsRole:nsRole` for Sun ONE directory, if groups are created with role inside Sun ONE
 - `memberOf:member` in Microsoft Active Directory Server
 - `groupMembership:member` for eDirectory
- Use the sample entries of `objectClass:attribute` pairs in the Group member ID map field.
 - `dominoGroup:member` for Lotus Domino
 - `groupOfNames:member` for eDirectory

Results

While using the direct method, dynamic groups, recursive groups, and static groups can be returned as multiple values of a single attribute. For example, in IBM Directory Server all group memberships, including the static groups, dynamic groups, and nested groups, can be returned using the `ibm-allGroups` attribute. In Sun ONE, all roles, including managed roles, filtered roles, and nested roles, are calculated

using the nsRole attribute. If an LDAP server can use the nsRole attribute, dynamic groups, nested groups, and static groups are all supported by WebSphere Application Server.

Some LDAP servers do not have recursive computing functionality. For example, although Microsoft Active Directory server has direct group search capability using the memberOf attribute, this attribute lists the groups beneath, which the group is directly nested only and does not contain the recursive list of nested predecessors. The Lotus Domino LDAP server only supports the indirect method to locate the group memberships for a user. You cannot obtain recursive group memberships from a Domino server directly. For LDAP servers without recursive searching capability, WebSphere Application Server security provides a recursive function that is enabled by clicking **Perform a Nested Group Search** in the Advanced LDAP user registry settings. Select this option only if your LDAP server does not provide recursive searches and you want a recursive search.

Configuring dynamic and nested group support for the SunONE or iPlanet Directory Server:

Configure dynamic and nested groups to simplify WebSphere Application Server security management and increase its effectiveness and flexibility.

Before you begin

To use dynamic and nested groups with WebSphere Application Server security, you must be running WebSphere Application Server Version 5.1.1 or later. Refer to “Dynamic groups and nested group support” on page 243 for more information on this topic.

1. In the administrative console for WebSphere Application Server, click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
3. Select SunONE for the type of LDAP server.
4. Select the **Ignore case for authorization** option.
5. Under Additional Properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
6. Change the Group filter setting to `&(cn=%v)(objectclass=ldapsubentry)`.
7. Change the Group member ID map setting to `nsRole:nsRole`.
8. Click **Apply** or **OK** to validate the changes.

Configuring dynamic and nested group support for the IBM Tivoli Directory Server:

Configure dynamic and nested groups to simplify WebSphere Application Server security management and increase its effectiveness and flexibility.

Before you begin

When creating groups, ensure that nested and dynamic group memberships work correctly.

1. In the administrative console for WebSphere Application Server, click **Security > Global security**.
2. Under User account repository, click **Standalone LDAP registry**, and click **Configure**.
3. Select IBM Tivoli Directory Server for the type of LDAP server.
4. Under Additional properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
5. Change the Group filter value to `(&(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames)(objectclass=groupOfURLs)))`.
6. Change the Group member ID map value to `ibm-allGroups:member;ibm-allGroups:uniqueMember`.
7. Click **Apply** or **OK** to validate the changes.

8. Verify that Auxiliary object class field on the Add an LDAP entry panel for your IBM Tivoli Directory server has the appropriate value. When you create a nested group, the Auxiliary object class value is `ibm-nestedGroup`. When you create a dynamic group, the Auxiliary object class value is `ibm-dynamicGroup`.

Configuring multiple LDAP servers for user registry failover

WebSphere Application Server security can be configured to attempt failovers between multiple Lightweight Directory Access Protocol (LDAP) hosts.

Before you begin

The multiple LDAP servers involved in the failover can be replicas that are replicated from the same master LDAP server, or they can be any LDAP host with the same schema. That is any LDAP host that contains data that is imported from the same LDAP data interchange format (LDIF) file.

Note: When WebSphere Application Server attempts failovers between multiple Lightweight Directory Access Protocol (LDAP) hosts, system properties are exchanged. WebSphere Application Server Version 6.1.0 manages the SSL configuration and these system properties. You cannot expect to set system properties yourself and expect the failover to succeed.

1. Start the deployment manager process.
 - a. Start the Command Prompt application.
 - b. Change directories to `profile_root\bin`.
 - c. Enter `startManager`.
2. Start the wsadmin Command Prompt application.
 - a. Start the Command Prompt application.
 - b. Change directories to `profile_root\bin`.
 - c. Enter the following command:

```
wsadmin -user username -password password
```
3. Configure a second LDAP server for failover.
 - a. Enter the following command to set the failover LDAP server hostname:

```
set ldapServer [ldap server hostname]
```
 - b. Enter the following command to set the LDAP server port number:

```
set ldapPort [ldap server port]
```
 - c. Enter the following command to set the WebSphere LDAP failover variable:

```
set Attrs2 [list [list hosts [list [list [list host $ldapServer] [list port $ldapPort]]]]]
```
 - d. Modify the LDAP configuration to add the failover LDAP server by entering the following command:

```
set result [$AdminConfig list LDAPUserRegistry]
```
 - e. Find the LDAP server configID by entering the following command:

```
$AdminConfig modify $result $Attrs2
```
 - f. Enter the following command to save the configuration change:

```
$AdminConfig save
```
 - g. Enter `exit` to quit the Command Prompt application. The following is an example of the Command Prompt application output:

```
wsadmin>set ldapServer [list xxxx.xxxx.xxx.com]
xxxx.xxxx.xxx.com
wsadmin>set ldapPort [list NNN]
NNN
wsadmin>set Attrs2 [list [list hosts [list [list [list host $ldapServer] [list port $ldapPort]]]]]
{hosts {{{host xxxx.xxxx.xxx.com} {port NNN}}}}
wsadmin> set result [$AdminConfig list LDAPUserRegistry]
```

```
(cells/Father2Cell01|security.xml#LDAPUserRegistry_1)
wsadmin>$AdminConfig modify $result $Attrs2
```

```
wsadmin>$AdminConfig save
```

4. Review the configuration change by opening the security.xml file with a text editor and review the new entry.
5. Stop the deployment manager.
 - a. Start the Command Prompt application.
 - b. Change directories to *profile_root*\bin.
 - c. To stop the deployment manager, enter the following command:

```
stopManager -user username -password password
```

Related tasks

“Testing an LDAP server for user registry failover”

After configuring a Lightweight Directory Access Protocol (LDAP) host for failover you should test the failover server by stopping the main LDAP server.

“Deleting LDAP endpoints using wsadmin” on page 140

You can delete Lightweight Directory Access Protocol (LDAP) endpoints for a user registry by using the WebSphere Application Server administrative tool (wsadmin).

Testing an LDAP server for user registry failover

After configuring a Lightweight Directory Access Protocol (LDAP) host for failover you should test the failover server by stopping the main LDAP server.

Before you begin

This task assumes the following setup:

- Deployment Manager is installed on the primary LDAP server running Application Server version 6.0.2 or higher.
 - All other LDAP hosts are Active Directory machines with similar user registry designs.
 - At least one of the other LDAP hosts has been configured for failover.
1. Stop the Active Directory Server on the failover server.
 2. Start the deployment manager process.
 - a. Start the Command Prompt application.
 - b. Change directories to *profile_root*\bin.
 - c. Enter startManager.
 3. Review the SystemOut.log file to see if the LDAP failover happened. The sample text is an example of a SystemOut.log file that records a successful failover:

```
[7/11/05 15:38:31:324 EDT] 0000000a LdapRegistryI A SECJ0418I:
Cannot connect to the LDAP server ldap://xxxx.xxxxx.xxxx.com:NNN. {primary LDAP server}
[7/11/05 15:38:32:486 EDT] 0000000a UserRegistryI A SECJ0136I:
Custom Registry:com.ibm.ws.security.registry.ldap.LdapRegistryImpl has been initialized
[7/11/05 15:38:53:787 EDT] 0000000a LdapRegistryI A SECJ0419I:
The user registry is currently connected to the LDAP server ldap://xxxx.xxxxx.xxxx.com:NNN. {failover LDAP server}
...
[7/11/05 15:39:35:667 EDT] 0000000a WsServerImpl A WSVR0001I: Server dmgr open for e-business
```
 4. Log into the console to see working and non-working cases.
 - a. Start a browser.
 - b. Browse to <http://localhost:9060/admin>.
 - c. Type in your user ID and password and click **OK**.
 - d. Log out of the Administrative Console.

- e. Type in DummyAdmin as the user ID and dummy1admin as your password and click **OK**. This should fail proving WebSphere Application Server is connected to the other LDAP server. Please make sure that on a production system the user registries are identical so this problem does not happen when switching between LDAP servers.
5. Stop the deployment manager.
 - a. Start the Command Prompt application.
 - b. Change directories to *profile_root\bin*.
 - c. To stop the deployment manager, enter the following command:


```
stopManager -user username -password password
```

Related tasks

“Configuring multiple LDAP servers for user registry failover” on page 138

WebSphere Application Server security can be configured to attempt failovers between multiple Lightweight Directory Access Protocol (LDAP) hosts.

Deleting LDAP endpoints using wsadmin

You can delete Lightweight Directory Access Protocol (LDAP) endpoints for a user registry by using the WebSphere Application Server administrative tool (wsadmin).

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Set the LDAP variable and display a list of LDAP endpoint objects. Enter the following commands:

Using Jacl:

```
set ldap [$AdminConfig list LDAPUserRegistry]
```

```
$AdminConfig list EndPoint $ldap
```

Using Jython:

```
ldap=AdminConfig.list["LDAPUserRegistry"]
```

```
print AdminConfig.show(ldap)
```

For the Jython language, you can obtain the endpoint from the host variable after running the previous command.

3. Display a list of LDAP endpoint objects. Enter the following command for each object:

Using Jacl:

```
$AdminConfig showall End_Point_Object
```

Using Jython:

```
AdminConfig.showall("End_Point_Object")
```

4. Delete an LDAP endpoint object. Enter the following command:

Using Jacl:

```
$AdminConfig remove End_Point_Object
```

Using Jython:

```
AdminConfig.remove ("End_Point_Object")
```

5. Save your configuration changes: Enter the following command:

Using Jacl:

```
$AdminConfig save
```

Using Jython:

```
AdminConfig.save()
```

Related tasks

“Testing an LDAP server for user registry failover” on page 139

After configuring a Lightweight Directory Access Protocol (LDAP) host for failover you should test the failover server by stopping the main LDAP server.

Updating LDAP binding information

Use this information to dynamically update security LDAP binding information by switching to a different binding identity.

About this task

You can dynamically update Lightweight Directory Access Protocol (LDAP) binding information without first stopping and restarting WebSphere Application Server by using the **wsadmin** tool.

The `resetLdapBindInfo` method in `SecurityAdmin` MBean is used to dynamically update LDAP binding information at WebSphere Application Server security runtime, and it takes the bind distinguished name (DN) and bind password parameters as input. The `resetLdapBindInfo` method validates the bind information against the LDAP server. If validation passes, new binding information is stored in `security.xml`, and a copy of the information is placed in WebSphere Application Server security runtime.

The MBean method also synchronizes the binding information change in `security.xml` from cell to nodes.

If the new binding information is `null`, `null`, the `resetLdapBindInfo` method first extracts LDAP binding information, including bind DN, bind password, and target binding host from WebSphere Application Server security configuration in `security.xml`. It then pushes the binding information to WebSphere Application Server security runtime.

There are two ways to dynamically update WebSphere Application Server security LDAP binding information using the `SecurityAdmin` MBean through `wsadmin`:

- “Switching to a different binding identity”
- “Switching to a failover LDAP host” on page 142

Switching to a different binding identity:

About this task

To dynamically update security LDAP binding information by switching to a different binding identity:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
3. Create a new bind DN. It must have the same access authority as the current bind DN.
4. Run the `SecurityAdmin` MBean across all of the processes (deployment manager, nodes, and application servers) to validate the new binding information, to save it to `security.xml`, and to push the new binding information to the runtime.

Example

The following is a sample Jacl file for step 4:

```
proc LDAPReBind {args} {
    global AdminConfig AdminControl ldapBindDn ldapBindPassword
    set ldapBindDn [lindex $args 0]
    set ldapBindPassword [lindex $args 1]
    set secMBeans [$AdminControl queryNames type=SecurityAdmin,*]
    set plist [list $ldapBindDn $ldapBindPassword]
```

```

    foreach secMBean $secMBeans {
        set result [$AdminControl invoke $secMBean resetLdapBindInfo $plist]
    }
}

```

Switching to a failover LDAP host: About this task

To dynamically update security LDAP binding information by switching to a failover LDAP host:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Standalone LDAP registry** and click **Configure**.
3. Change the password for bind DN on one LDAP server (it can be the primary or the backup).
4. Update the new bind DN password to WebSphere Application security runtime by calling `resetLdapBindInfo` with the bind DN and by using its new password as a parameter.
5. Use the new bind DN password for all of the other LDAP servers. The binding information is now consistent across WebSphere Application Server and the LDAP servers.

Note: If you call `resetLdapBindInfo` with `null`, `null` as input parameters, WebSphere Application Server security runtime completes the following steps:

- a. Reads the bind DN, bind password, and target LDAP hosts from `security.xml`.
- b. Refreshes the cached connection to the LDAP server.

If you configure security to use multiple LDAP servers, this MBean call forces WebSphere Application Server security to reconnect to the first available LDAP host in the list. For example, if three LDAP servers are configured in the order of L1, L2, and L3, the reconnection process always starts with the L1 server.

When LDAP failover is configured by associating a single hostname to multiple IP addresses, entering an invalid password can cause multiple LDAP bind retries. With the default settings, the number of LDAP bind retries is equal to one more than the number of associated IP addresses. This means a single invalid login attempt can cause the LDAP account to be locked. If the `com.ibm.websphere.security.ldap.retryBind` custom property is set to `false`, LDAP bind calls are not retried.

Configuring standalone custom registries

Use the following information to configure standalone custom registries through the administrative console.

Before you begin

Before you begin this task, implement and build the `UserRegistry` interface. For more information on developing standalone custom registries refer to “Developing standalone custom registries” on page 731. The following steps are required to configure standalone custom registries through the administrative console.

1. Click **Security > Global security**.
2. Under User account repositories, select **Standalone custom registry** and click **Configure**.
3. Enter a valid user name in the Primary administrative user name field. This ID is the security server ID, which is only used for WebSphere Application Server security and is not associated with the system process that runs the server. The server calls the local operating system registry to authenticate and obtain privilege information about users by calling the native APIs in that particular registry.
4. Enter the complete location of the dot-separated class name that implements the `com.ibm.websphere.security.UserRegistry` interface in the Custom registry class name field. For the sample, this file name is `com.ibm.websphere.security.FileRegistrySample`.

The file exists in the WebSphere Application Server class path preferably in the *app_server_root/lib/ext* directory.

This file exists in all the product processes. Thus, this file exists in the cell class path and in all of the node class paths.

Note: The sample provided is intended to familiarize you with this feature. Do not use this sample in an actual production environment.

5. Add your custom registry class name to the class path. It is recommended that you add the Java Archive (JAR) file that contains your custom user registry implementation to the *app_server_root/classes* directory.
6. Optional: Select the **Ignore case for authorization** option for the authorization to perform a case insensitive check. Enabling this option is necessary only when your user registry is case insensitive and does not provide a consistent case when queried for users and groups.
7. Click **Apply** if you have any other additional properties to enter for the registry initialization.
8. Optional: Enter additional properties to initialize your implementation.
 - a. Click **Custom properties > New**.
 - b. Enter the property name and value.

For the sample, enter the following two properties. It is assumed that the *users.props* file and the *groups.props* file are in the *customer_sample* directory under the product installation directory. You can place these properties in any directory that you choose and reference their locations through custom properties. However, make sure that the directory has the appropriate access permissions.

Property name	Property value
usersFile	<code>\${USER_INSTALL_ROOT}/customer_sample /users.props</code>
groupsFile	<code>\${USER_INSTALL_ROOT}/customer_sample /groups.props</code>

Samples of these two properties are available in “users.props file” on page 164 and “groups.props file” on page 164.

The **Description**, **Required**, and **Validation Expression** fields are not used and can remain blank.

In a Network Deployment environment where multiple WebSphere Application Server processes exist, such as cell and multiple nodes in different machines, these properties are available for each process. Use the *USER_INSTALL_ROOT* relative name to locate any files, as this name expands to the product installation directory. If this name is not used, ensure that the files exist in the same location in all the nodes.

WebSphere Application Server version 4-based custom user registry is migrated to the custom user registry based on the `com.ibm.websphere.security.UserRegistry` interface.

- c. Click **Apply**.
 - d. Repeat this step to add other additional properties.
9. Click **Security > Global security**.
10. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone custom registry**, and click **Configure**.
11. Select either the **Automatically generated server identity** or **Server identity that is stored in the repository** option. If you select the **Server identity that is stored in the repository** option, enter the following information:

Server user ID or administrative user on a Version 6.0.x node

Specify the short name of the account that is chosen in the second step.

Server user password

Specify the password of the account that is chosen in the second step.

12. Click **OK** and complete the required steps to turn on security.

Results

This set of steps is required to set up the standalone custom registry and to enable security in WebSphere Application Server.

Note: The security component of WebSphere Application Server expands a selected list of variables when enabling security. See Variable settings for more detail.

What to do next

1. Complete the remaining steps, if you are enabling security.
2. Validate the user and password. Save and synchronize in the cell environment.
3. After security is turned on, save, stop, and start all the product servers, including cell, nodes, and all of the application servers, for any changes to take effect. If the server comes up without any problems, the setup is correct.

Standalone custom registry settings

Use this page to configure the standalone custom registry.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone custom registry**, and click **Configure**.

After the properties are set in this panel, click **Apply**. Under Additional Properties, click **Custom properties** to include additional properties that the custom user registry requires.

Note: Custom properties might include information such as specifying lists of users or groups.

When security is enabled and any of these custom user registry settings change, go to the Global security panel and click **Apply** to validate the changes.

WebSphere Application Server Version 6.1 distinguishes between the user identities for administrators who manage the environment and server identities for authenticating server to server communications. In most cases, server identities are automatically generated and are not stored in a repository.

However, if you are adding a Version 5.0.x or 6.0.x node to a Version 6.1 cell, you must ensure that the Version 5.x or Version 6.0.x server identity and password are defined in the repository for this cell. Enter the server user identity and password on this panel.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your custom user registry.

The user name is used to log onto the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Automatically generated server identity:

Enables the application server to generate the server identity, which is recommended for environments that contain only Version 6.1 or later nodes. Automatically generated server identities are not stored in a user repository.

You can change this server identity on the Authentication mechanisms and expiration panel. To access the Authentication mechanisms and expiration panel, click **Security > Global security > Authentication mechanisms and expiration**. Change the value of the Internal server ID field.

Default: Enabled

Server identity that is stored in the repository:

Specifies a user identity in the repository that is used for internal process communication. Cells that contain Version 5.x or 6.0.x nodes require a server user identity that is defined in the active user repository.

Default: Enabled

Server user ID or administrative user on a Version 6.0.x node:

Specifies the user ID that is used to run the application server for security purposes.

Password:

Specifies the password that corresponds to the server ID.

Custom registry class name:

Specifies a dot-separated class name that implements the com.ibm.websphere.security.UserRegistry interface.

Put the custom registry class name in the class path. A suggested location is the following directory.

- `%install_root%/lib/ext`

Data type: String
Default: com.ibm.websphere.security.FileRegistrySample

Ignore case for authorization:

Indicates that a case-insensitive authorization check is performed when you use the default authorization.

Default: Disabled
Range: Enabled or Disabled

Standalone custom registry wizard settings

A wizard page exists in the administrative console to aid in viewing the basic settings necessary to connect the application server to an existing standalone custom registry. After you have viewed the basic settings, you can also modify the existing standalone customer registry configuration using the administrative console.

To view this security wizard page, complete the following steps:

1. Click **Security > Global security > Security configuration wizard**.
2. Select your protection settings and click **Next**.

3. Select the **Standalone custom registry** option and click **Next**.

You can modify your standalone custom registry configuration by completing the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone custom registry**, and click **Configure**.
3. Enter additional properties to initialize your implementation
 - Click **Custom properties > New**.
 - Enter the property name and value. For the sample, enter the following two properties. It is assumed that the users.props file and the groups.props file are in the *customer_sample* directory under the product installation directory. You can place these properties in any directory that you choose and reference their locations through Custom properties. However, make sure that the directory has the appropriate access permissions.

Property name	Property value
usersFile	<code>\${USER_INSTALL_ROOT}/customer_sample /users.props</code>
groupsFile	<code>\${USER_INSTALL_ROOT}/customer_sample /groups.props</code>

Samples of these two properties are available in reference topics for the users.props file and the groups.props file. See the related links below for more information.

The Description, Required, and Validation Expression fields are not used and can remain blank.

WebSphere Application Server Version 4 based custom user registry is migrated to the custom user registry based on the com.ibm.websphere.security.UserRegistry interface.

- Click **Apply**.

Primary administrative user name:

Specifies the name of a user with administrative privileges that is defined in your custom user registry.

The user name is used to log onto the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Custom registry class name:

Specifies a dot-separated class name that implements the com.ibm.websphere.security.UserRegistry interface.

Put the custom registry class name in the class path. A suggested location is the following directory.

- `%install_root%/lib/ext`

Data type:

String

Default:

com.ibm.websphere.security.FileRegistrySample

Ignore case for authorization:

Indicates that a case-insensitive authorization check is performed when you use the default authorization.

Default: Disabled
Range: Enabled or Disabled

FileRegistrySample.java file

This provides an example of the FileRegistrySample.java file.

The user and group information required by this sample is contained in the "users.props file" on page 164 and "groups.props file" on page 164 files.

Note: The samples that are provided are intended to familiarize you with this feature. Do not use these samples in an actual production environment.

The contents of the FileRegistrySample.java file:

```
//
// 5639-D57, 5630-A36, 5630-A37, 5724-D18
// (C) COPYRIGHT International Business Machines Corp. 1997, 2005
// All Rights Reserved * Licensed Materials - Property of IBM
////-----
// This program may be used, run, copied, modified and distributed
// without royalty for the purpose of developing, using, marketing, or
// distributing.
//-----
//

// This sample is for the custom user registry feature in WebSphere
// Application Server.

import java.util.*;
import java.io.*;
import java.security.cert.X509Certificate;
import com.ibm.websphere.security.*;
/**
 * The main purpose of this sample is to demonstrate the use of the
 * custom user registry feature available in WebSphere Application Server. This
 * sample is a file-based registry sample where the users and the groups
 * information is listed in files (users.props and groups.props). As such
 * simplicity and not the performance was a major factor. This
 * sample should be used only to get familiarized with this feature. An
 * actual implementation of a realistic registry should consider various
 * factors like performance, scalability, thread safety, and so on.
 */
public class FileRegistrySample implements UserRegistry {

    private static String USERFILENAME = null;
    private static String GROUPFILENAME = null;

    /** Default Constructor */
    public FileRegistrySample() throws java.rmi.RemoteException {
    }

    /**
     * Initializes the registry. This method is called when creating the
     * registry.
     *
     * @param props - The registry-specific properties with which to
     *                initialize the custom registry
     * @exception CustomRegistryException
     *                if there is any registry-specific problem
     */
    public void initialize(java.util.Properties props)
```

```

    throws CustomRegistryException {
try {
    /* try getting the USERFILENAME and the GROUPFILENAME from
    * properties that are passed in (For example, from the
    * administrative console). Set these values in the administrative
    * console. Go to the special custom settings in the custom
    * user registry section of the Authentication panel.
    * For example:
    * usersFile c:/temp/users.props
    * groupsFile c:/temp/groups.props
    */
    if (props != null) {
        USERFILENAME = props.getProperty("usersFile");
        GROUPFILENAME = props.getProperty("groupsFile");
    }

} catch(Exception ex) {
    throw new CustomRegistryException(ex.getMessage(),ex);
}

if (USERFILENAME == null || GROUPFILENAME == null) {
    throw new CustomRegistryException("users/groups information missing");
}

} /**
* Checks the password of the user. This method is called to authenticate
* a user when the user's name and password are given.
*
* @param userSecurityName the name of user
* @param password the password of the user
* @return a valid userSecurityName. Normally this is
*         the name of same user whose password was checked
*         but if the implementation wants to return any other
*         valid userSecurityName in the registry it can do so
* @exception CheckPasswordFailedException if userSecurityName/
*         password combination does not exist
*         in the registry
* @exception CustomRegistryException if there is any registry-
*         specific problem
**/
public String checkPassword(String userSecurityName, String passwd)
    throws PasswordCheckFailedException,
        CustomRegistryException {
    String s,userName = null;
    BufferedReader in = null;

    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                int index1 = s.indexOf(":",index+1);
                // check if the userSecurityName:passwd combination exists
                if ((s.substring(0,index)).equals(userSecurityName) &&
                    s.substring(index+1,index1).equals(passwd)) {
                    // Authentication successful, return the userID.
                    userName = userSecurityName;
                    break;
                }
            }
        }
    }
    } catch(Exception ex) {

```

```

        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    if (userName == null) {
        throw new PasswordCheckFailedException("Password check failed for user: "
            + userSecurityName);
    }

    return userName;
} /**
 * Maps an X.509 format certificate to a valid user in the registry.
 * This is used to map the name in the certificate supplied by a browser
 * to a valid userSecurityName in the registry
 *
 * @param cert the X509 certificate chain
 * @return The mapped name of the user userSecurityName
 * @exception CertificateMapNotSupportedException if the
 * particular certificate is not supported.
 * @exception CertificateMapFailedException if the mapping of
 * the certificate fails.
 * @exception CustomRegistryException if there is any registry
 * -specific problem
 */
public String mapCertificate(X509Certificate[] cert)
    throws CertificateMapNotSupportedException,
        CertificateMapFailedException,
        CustomRegistryException {
    String name=null;
    X509Certificate cert1 = cert[0];
    try {
        // map the SubjectDN in the certificate to a userID.
        name = cert1.getSubjectDN().getName();
    } catch(Exception ex) {
        throw new CertificateMapNotSupportedException(ex.getMessage(),ex);
    }

    if(!isValidUser(name)) {
        throw new CertificateMapFailedException("user: " + name
            + " is not valid");
    }
    return name;
} /**
 * Returns the realm of the registry.
 *
 * @return the realm. The realm is a registry-specific string
 * indicating the realm or domain for which this registry
 * applies. For example, for OS/400 or AIX this would be
 * the host name of the system whose user registry this
 * object represents. If null is returned by this method,
 * realm defaults to the value of "customRealm". It is
 * recommended that you use your own value for realm.
 *
 * @exception CustomRegistryException if there is any registry-
 * specific problem
 */
public String getRealm()
    throws CustomRegistryException {
    String name = "customRealm";
    return name;
} /**

```

```

* Gets a list of users that match a pattern in the registry.
* The maximum number of users returned is defined by the limit
* argument.
* This method is called by the administrative console and scripting
* (command line) to make the users in the registry available for
* adding them (users) to roles.
*
* @param      pattern the pattern to match. (For example, a* will
* match all userSecurityNames starting with a)
* @param      limit the maximum number of users that should be
* returned. This is very useful in situations where
* there are thousands of users in the registry and
* getting all of them at once is not practical. The
* default is 100. A value of 0 implies get all the
* users and hence must be used with care.
* @return     a Result object that contains the list of users
* requested and a flag to indicate if more users
* exist.
* @exception  CustomRegistryException if there is any registry-
* specific problem
**/
public Result getUsers(String pattern, int limit)
    throws CustomRegistryException {
    String s;
    BufferedReader in = null;
    List allUsers = new ArrayList();
    Result result = new Result();
    int count = 0;
    int newLimit = limit+1;
    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                String user = s.substring(0,index);
                if (match(user,pattern)) {
                    allUsers.add(user);
                    if (limit !=0 && ++count == newLimit) {
                        allUsers.remove(user);
                        result.setHasMore();
                        break;
                    }
                }
            }
        }
    } catch (Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    result.setList(allUsers);
    return result;
} /**
* Returns the display name for the user specified by
* userSecurityName.
*
* This method may be called only when the user information
* is displayed (information purposes only, for example, in
* the administrative console) and hence not used in the actual
* authentication or authorization purposes. If there are no
* display names in the registry return null or empty string.

```

```

*
* In WebSphere Application Server 4.x custom registry, if you
* had a display name for the user and if it was different from the
* security name, the display name was returned for the EJB
* methods getCallerPrincipal() and the servlet methods
* getUserPrincipal() and getRemoteUser().
* In WebSphere Application Server Version 5.x and later, for the
* same methods, the security name will be returned by default.
* This is the recommended way as the display name is not unique
* and might create security holes. However, for backward
* compatibility if you need the display name to be returned
* set the property WAS_UseDisplayName to true.
*
*See the Information Center documentation for more information.
*
* @param      userSecurityName the name of the user.
* @return     the display name for the user. The display
*             name is a registry-specific string that
*             represents a descriptive, not necessarily
*             unique, name for a user. If a display name
*             does not exist return null or empty string.
* @exception  EntryNotFoundException if userSecurityName
*             does not exist.
* @exception  CustomRegistryException if there is any registry-
*             specific problem
**/
public String getUserDisplayName(String userSecurityName)
    throws CustomRegistryException,
           EntryNotFoundException {

    String s,displayName = null;
    BufferedReader in = null;

    if(!isValidUser(userSecurityName)) {
        EntryNotFoundException nsee = new EntryNotFoundException("user: "
            + userSecurityName + " is not valid");
        throw nsee;
    }

    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                int index1 = s.lastIndexOf(":");
                if ((s.substring(0,index)).equals(userSecurityName)) {
                    displayName = s.substring(index1+1);
                    break;
                }
            }
        }
    } catch(Exception ex) {
        throw new CustomRegistryException(ex.getMessage(), ex);
    } finally {
        fileClose(in);
    }

    return displayName;
}

/**
* Returns the unique ID for a userSecurityName. This method is called

```



```

* when creating a credential for a user.
*
* @param    userSecurityName - The name of the user.
* @return   The unique ID of the user. The unique ID for a user
*           is the stringified form of some unique, registry-specific,
*           data that serves to represent the user. For example, for
*           the UNIX user registry, the unique ID for a user can be
*           the UID.
* @exception EntryNotFoundException if userSecurityName does not
*           exist.
* @exception CustomRegistryException if there is any registry-
*           specific problem
**/
public String getUniqueId(String userSecurityName)
    throws CustomRegistryException,
           EntryNotFoundException {

    String s,uniqueUsrId = null;
    BufferedReader in = null;
    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                int index1 = s.indexOf(":", index+1);
                if ((s.substring(0,index)).equals(userSecurityName)) {
                    int index2 = s.indexOf(":", index1+1);
                    uniqueUsrId = s.substring(index1+1,index2);
                    break;
                }
            }
        }
    } catch(Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    if (uniqueUsrId == null) {
        EntryNotFoundException nsee =
            new EntryNotFoundException("Cannot obtain uniqueId for user: "
                + userSecurityName);
        throw nsee;
    }

    return uniqueUsrId;
} /**
* Returns the name for a user given its unique ID.
*
* @param    uniqueUserId - The unique ID of the user.
* @return   The userSecurityName of the user.
* @exception EntryNotFoundException if the unique user ID does not exist.
* @exception CustomRegistryException if there is any registry-specific
*           problem
**/
public String getUserSecurityName(String uniqueUserId)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s,usrSecName = null;
    BufferedReader in = null;
    try {
        in = fileOpen(USERFILENAME);

```

```

while ((s=in.readLine())!=null)
{
    if (!(s.startsWith("#") || s.trim().length() <=0 )) {
        int index = s.indexOf(":");
        int index1 = s.indexOf(":", index+1);
        int index2 = s.indexOf(":", index1+1);
        if ((s.substring(index1+1,index2)).equals(uniqueUserId)) {
            usrSecName = s.substring(0,index);
            break;
        }
    }
}
} catch (Exception ex) {
    throw new CustomRegistryException(ex.getMessage(), ex);
} finally {
    fileClose(in);
}

if (usrSecName == null) {
    EntryNotFoundException ex =
        new EntryNotFoundException("Cannot obtain the
            user securityName for " + uniqueUserId);
    throw ex;
}

return usrSecName;

} /**
 * Determines if the userSecurityName exists in the registry
 *
 * @param    userSecurityName - The name of the user
 * @return    True if the user is valid; otherwise false
 * @exception CustomRegistryException if there is any registry-
 *           specific problem
 * @exception RemoteException as this extends java.rmi.Remote
 *           interface
 **/
public boolean isValidUser(String userSecurityName)
    throws CustomRegistryException {
    String s;
    boolean isValid = false;
    BufferedReader in = null;
    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                if ((s.substring(0,index)).equals(userSecurityName)) {
                    isValid=true;
                    break;
                }
            }
        }
    }
} catch (Exception ex) {
    throw new CustomRegistryException(ex.getMessage(), ex);
} finally {
    fileClose(in);
}

return isValid;
}
/**

```

```

* Gets a list of groups that match a pattern in the registry
* The maximum number of groups returned is defined by the
* limit argument. This method is called by administrative console
* and scripting (command line) to make available the groups in
* the registry for adding them (groups) to roles.
*
* @param      pattern the pattern to match. (For example, a* matches
*              all groupSecurityNames starting with a)
* @param      limit Limits the maximum number of groups to return
*              This is very useful in situations where there
*              are thousands of groups in the registry and getting all
*              of them at once is not practical. The default is 100.
*              A value of 0 implies get all the groups and hence must
*              be used with care.
* @return     A Result object that contains the list of groups
*              requested and a flag to indicate if more groups exist.
* @exception  CustomRegistryException if there is any registry-specific
*              problem
**/
public Result getGroups(String pattern, int limit)
    throws CustomRegistryException {
    String s;
    BufferedReader in = null;
    List allGroups = new ArrayList();      Result result = new Result();
    int count = 0;
    int newLimit = limit+1;
    try {
        in = fileOpen(GROUPFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                String group = s.substring(0,index);
                if (match(group,pattern)) {
                    allGroups.add(group);
                    if (limit !=0 && ++count == newLimit) {
                        allGroups.remove(group);
                        result.setHasMore();
                        break;
                    }
                }
            }
        }
    } catch (Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    result.setList(allGroups);
    return result;
}

/**
* Returns the display name for the group specified by groupSecurityName.
* For this version of WebSphere Application Server, the only usage of
* this method is by the clients (administrative console and scripting)
* to present a descriptive name of the user if it exists.
*
* @param groupSecurityName the name of the group.
* @return the display name for the group. The display name
*         is a registry-specific string that represents a
*         descriptive, not necessarily unique, name for a group.

```

```

*         If a display name does not exist return null or empty
*         string.
* @exception EntryNotFoundException if groupSecurityName does
*         not exist.
* @exception CustomRegistryException if there is any registry-
*         specific problem
**/
public String getGroupDisplayName(String groupSecurityName)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s,displayName = null;
    BufferedReader in = null;

    if(!isValidGroup(groupSecurityName)) {
        EntryNotFoundException nsee = new EntryNotFoundException("group: "
            + groupSecurityName + " is not valid");
        throw nsee;
    }

    try {
        in = fileOpen(GROUPFILENAME);
        while ((s=in.readLine())!=null)
            {
                if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                    int index = s.indexOf(":");
                    int index1 = s.lastIndexOf(":");
                    if ((s.substring(0,index)).equals(groupSecurityName)) {
                        displayName = s.substring(index1+1);
                        break;
                    }
                }
            }
    } catch(Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    return displayName;
}

/**
* Returns the Unique ID for a group.

* @param    groupSecurityName the name of the group.
* @return   The unique ID of the group. The unique ID for
*           a group is the stringified form of some unique,
*           registry-specific, data that serves to represent
*           the group. For example, for the UNIX user registry,
*           the unique ID might be the GID.
* @exception EntryNotFoundException if groupSecurityName does
*         not exist.
* @exception CustomRegistryException if there is any registry-
*         specific problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getUniqueGroupId(String groupSecurityName)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s,uniqueGrpId = null;
    BufferedReader in = null;
    try {
        in = fileOpen(GROUPFILENAME);

```

```

while ((s=in.readLine())!=null)
{
    if (!(s.startsWith("#") || s.trim().length() <=0 )) {
        int index = s.indexOf(":");
        int index1 = s.indexOf(":", index+1);
        if ((s.substring(0,index)).equals(groupSecurityName)) {
            uniqueGrpId = s.substring(index+1,index1);
            break;
        }
    }
}
} catch(Exception ex) {
    throw new CustomRegistryException(ex.getMessage(),ex);
} finally {
    fileClose(in);
}

if (uniqueGrpId == null) {
    EntryNotFoundException nsee =
    new EntryNotFoundException("Cannot obtain the uniqueId for group: "
    + groupSecurityName);
    throw nsee;
}

return uniqueGrpId;
}

/**
 * Returns the Unique IDs for all the groups that contain the unique ID
 * of a user. Called during creation of a user's credential.
 *
 * @param    uniqueUserId the unique ID of the user.
 * @return   A list of all the group unique IDs that the unique user
 *           ID belongs to. The unique ID for an entry is the
 *           stringified form of some unique, registry-specific, data
 *           that serves to represent the entry. For example, for the
 *           UNIX user registry, the unique ID for a group might be
 *           the GID and the Unique ID for the user might be the UID.
 * @exception EntryNotFoundException if uniqueUserId does not exist.
 * @exception CustomRegistryException if there is any registry-specific
 *           problem
 */
public List getUniqueGroupIds(String uniqueUserId)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s,uniqueGrpId = null;
    BufferedReader in = null;
    List uniqueGrpIds=new ArrayList();
    try {
        in = fileOpen(USERFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                int index1 = s.indexOf(":", index+1);
                int index2 = s.indexOf(":", index1+1);
                if ((s.substring(index1+1,index2)).equals(uniqueUserId)) {
                    int lastIndex = s.lastIndexOf(":");
                    String subs = s.substring(index2+1,lastIndex);
                    StringTokenizer st1 = new StringTokenizer(subs, ",");
                    while (st1.hasMoreTokens())
                        uniqueGrpIds.add(st1.nextToken());
                    break;
                }
            }
        }
    }
}

```

```

        }
    }
} catch(Exception ex) {
    throw new CustomRegistryException(ex.getMessage(),ex);
} finally {
    fileClose(in);
}

return uniqueGrpIds;
}

/**
 * Returns the name for a group given its unique ID.
 *
 * @param    uniqueGroupId the unique ID of the group.
 * @return   The name of the group.
 * @exception EntryNotFoundException if the uniqueGroupId does
 *         not exist.
 * @exception CustomRegistryException if there is any registry-
 *         specific problem
 */
public String getGroupSecurityName(String uniqueGroupId)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s,grpSecName = null;
    BufferedReader in = null;
    try {
        in = fileOpen(GROUPFILENAME);
        while ((s=in.readLine())!=null)
        {
            if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                int index = s.indexOf(":");
                int index1 = s.indexOf(":", index+1);
                if ((s.substring(index+1,index1)).equals(uniqueGroupId)) {
                    grpSecName = s.substring(0,index);
                    break;
                }
            }
        }
    }
    catch (Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    }
    finally {
        fileClose(in);
    }

    if (grpSecName == null) {
        EntryNotFoundException ex =
            new EntryNotFoundException("Cannot obtain the group
            security name for: " + uniqueGroupId);
        throw ex;
    }

    return grpSecName;
}

/**
 * Determines if the groupSecurityName exists in the registry
 *
 * @param    groupSecurityName the name of the group
 * @return   True if the groups exists; otherwise false
 * @exception CustomRegistryException if there is any registry-

```

```

*           specific problem
**/
public boolean isValidGroup(String groupSecurityName)
    throws CustomRegistryException {
    String s;
    boolean isValid = false;
    BufferedReader in = null;
    try {
        in = fileOpen(GROUPFILENAME);
        while ((s=in.readLine())!=null)
            {
                if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                    int index = s.indexOf(":");
                    if ((s.substring(0,index)).equals(groupSecurityName)) {
                        isValid=true;
                        break;
                    }
                }
            }
    } catch (Exception ex) {
        throw new CustomRegistryException(ex.getMessage(),ex);
    } finally {
        fileClose(in);
    }

    return isValid;
}

/**
 * Returns the securityNames of all the groups that contain the user
 *
 * This method is called by the administrative console and scripting
 * (command line) to verify that the user entered for RunAsRole mapping
 * belongs to that role in the roles to user mapping. Initially, the
 * check is done to see if the role contains the user. If the role does
 * not contain the user explicitly, this method is called to get the groups
 * that this user belongs to so that a check can be made on the groups that
 * the role contains.
 *
 * @param    userSecurityName the name of the user
 * @return    A list of all the group securityNames that the user
 *            belongs to.
 * @exception EntryNotFoundException if user does not exist.
 * @exception CustomRegistryException if there is any registry-
 *            specific problem
 * @exception RemoteException as this extends the java.rmi.Remote
 *            interface
 **/
public List getGroupsForUser(String userName)
    throws CustomRegistryException,
           EntryNotFoundException {
    String s;
    List grpsForUser = new ArrayList();
    BufferedReader in = null;
    try {
        in = fileOpen(GROUPFILENAME);
        while ((s=in.readLine())!=null)
            {
                if (!(s.startsWith("#") || s.trim().length() <=0 )) {
                    StringTokenizer st = new StringTokenizer(s, ":");
                    for (int i=0; i<2; i++)
                        st.nextToken();
                    String subs = st.nextToken();
                }
            }
    }
}

```



```

        StringTokenizer st1 = new StringTokenizer(subs, ",");
        while (st1.hasMoreTokens()) {
            if((st1.nextToken().equals(userName)) {
                int index = s.indexOf(":");
                grpsForUser.add(s.substring(0,index));
            }
        }
    }
} catch (Exception ex) {
    if (!isValidUser(userName)) {
        throw new EntryNotFoundException("user: " + userName
            + " is not valid");
    }
    throw new CustomRegistryException(ex.getMessage(),ex);
} finally {
    fileClose(in);
}

return grpsForUser;
}

/**
 * Gets a list of users in a group.
 *
 * The maximum number of users returned is defined by the
 * limit argument.
 *
 * This method is being used by the WebSphere Application
 * Server Enterprise process choreographer (Enterprise) when
 * staff assignments are modeled using groups.
 *
 * In rare situations, if you are working with a registry where
 * getting all the users from any of your groups is not practical
 * (for example if there are a large number of users) you can create
 * the NotImplementedException for that particular group. Make sure
 * that if the process choreographer is installed (or if installed later)
 * the staff assignments are not modeled using these particular groups.
 * If there is no concern about returning the users from groups
 * in the registry it is recommended that this method be implemented
 * without creating the NotImplementedException.
 * @param groupSecurityName the name of the group
 * @param limit Limits the maximum number of users that should be
 * returned. This is very useful in situations where there
 * are lots of users in the registry and getting all of
 * them at once is not practical. A value of 0 implies
 * get all the users and hence must be used with care.
 * @return A Result object that contains the list of users
 * requested and a flag to indicate if more users exist.
 * @deprecated This method will be deprecated in future.
 * @exception NotImplementedException create this exception in rare
 * situations if it is not practical to get this information
 * for any of the group or groups from the registry.
 * @exception EntryNotFoundException if the group does not exist in
 * the registry
 * @exception CustomRegistryException if there is any registry-specific
 * problem
 */
public Result getUsersForGroup(String groupSecurityName, int limit)
    throws NotImplementedException,
        EntryNotFoundException,
        CustomRegistryException {
    String s, user;

```

```

BufferedReader in = null;
List usrsForGroup = new ArrayList();
int count = 0;
int newLimit = limit+1;
Result result = new Result();

try {
    in = fileOpen(GROUPFILENAME);
    while ((s=in.readLine())!=null)
    {
        if (!(s.startsWith("#") || s.trim().length() <=0 )) {
            int index = s.indexOf(":");
            if ((s.substring(0,index)).equals(groupSecurityName))
            {
                StringTokenizer st = new StringTokenizer(s, ":");
                for (int i=0; i<2; i++)
                    st.nextToken();
                String subs = st.nextToken();
                StringTokenizer st1 = new StringTokenizer(subs, ",");
                while (st1.hasMoreTokens()) {
                    user = st1.nextToken();
                    usrsForGroup.add(user);
                    if (limit !=0 && ++count == newLimit) {
                        usrsForGroup.remove(user);
                        result.setHasMore();
                        break;
                    }
                }
            }
        }
    }
} catch (Exception ex) {
    if (!isValidGroup(groupSecurityName)) {
        throw new EntryNotFoundException("group: "
            + groupSecurityName
            + " is not valid");
    }
    throw new CustomRegistryException(ex.getMessage(),ex);
} finally {
    fileClose(in);
}

result.setList(usrsForGroup);
return result;
}

/**
 * This method is implemented internally by the WebSphere Application
 * Server code in this release. This method is not called for the custom
 * registry implementations for this release. Return null in the
 * implementation.
 */
public com.ibm.websphere.security.cred.WSCredential
    createCredential(String userSecurityName)
        throws CustomRegistryException,
            NotImplementedException,
            EntryNotFoundException {

    // This method is not called.
    return null;
}

```

```

// private methods
private BufferedReader fileOpen(String fileName)
    throws FileNotFoundException {
    try {
        return new BufferedReader(new FileReader(fileName));
    } catch(FileNotFoundException e) {
        throw e;
    }
}

private void fileClose(BufferedReader in) {
    try {
        if (in != null) in.close();
    } catch(Exception e) {
        System.out.println("Error closing file" + e);
    }
}

private boolean match(String name, String pattern) {
    RegExpSample regexp = new RegExpSample(pattern);
    boolean matches = false;
    if(regexp.match(name))
        matches = true;
    return matches;
}
}

//-----
// The program provides the Regular Expression implementation
// used in the sample for the custom user registry (FileRegistrySample).
// The pattern matching in the sample uses this program to search for the
// pattern (for users and groups).
//-----

class RegExpSample
{
    private boolean match(String s, int i, int j, int k)
    {
        for(; k < expr.length; k++)
label10:
        {
            Object obj = expr[k];
            if(obj == STAR)
            {
                if(++k >= expr.length)
                    return true;
                if(expr[k] instanceof String)
                {
                    String s1 = (String)expr[k++];
                    int l = s1.length();
                    for(; (i = s.indexOf(s1, i)) >= 0; i++)
                        if(match(s, i + 1, j, k))
                            return true;
                }
                return false;
            }
            for(; i < j; i++)
                if(match(s, i, j, k))
                    return true;
        }
        return false;
    }
}

```

```

    }
    if(obj == ANY)
    {
        if(++i > j)
            return false;
        break label0;
    }
    if(obj instanceof char[][])
    {
        if(i >= j)
            return false;
        char c = s.charAt(i++);
        char ac[][] = (char[][])obj;
        if(ac[0] == NOT)
        {
            for(int j1 = 1; j1 < ac.length; j1++)
                if(ac[j1][0] <= c && c <= ac[j1][1])
                    return false;

            break label0;
        }
        for(int k1 = 0; k1 < ac.length; k1++)
            if(ac[k1][0] <= c && c <= ac[k1][1])
                break label0;

        return false;
    }
    if(obj instanceof String)
    {
        String s2 = (String)obj;
        int i1 = s2.length();
        if(!s.regionMatches(i, s2, 0, i1))
            return false;
        i += i1;
    }
}

return i == j;
}

public boolean match(String s)
{
    return match(s, 0, s.length(), 0);
}

public boolean match(String s, int i, int j)
{
    return match(s, i, j, 0);
}

public RegExpSample(String s)
{
    Vector vector = new Vector();
    int i = s.length();
    StringBuffer stringbuffer = null;
    Object obj = null;
    for(int j = 0; j < i; j++)
    {
        char c = s.charAt(j);
        switch(c)
        {
            case 63: /* '?' */
                obj = ANY;

```

```

        break;

case 42: /* '*' */
    obj = STAR;
    break;

case 91: /* '[' */
    int k = ++j;
    Vector vector1 = new Vector();
    for(; j < i; j++)
    {
        c = s.charAt(j);
        if(j == k && c == '^')
        {
            vector1.addElement(NOT);
            continue;
        }
        if(c == '\\')
        {
            if(j + 1 < i)
                c = s.charAt(++j);
        }
        else
            if(c == ']')
                break;
        char c1 = c;
        if(j + 2 < i && s.charAt(j + 1) == '-')
            c1 = s.charAt(j += 2);
        char ac1[] = {
            c, c1
        };
        vector1.addElement(ac1);
    }

    char ac[][] = new char[vector1.size()][2];
    vector1.copyInto(ac);
    obj = ac;
    break;

case 92: /* '\\' */
    if(j + 1 < i)
        c = s.charAt(++j);
    break;

}
if(obj != null)
{
    if(stringbuffer != null)
    {
        vector.addElement(stringbuffer.toString());
        stringbuffer = null;
    }
    vector.addElement(obj);
    obj = null;
}
else
{
    if(stringbuffer == null)
        stringbuffer = new StringBuffer();
    stringbuffer.append(c);
}
}

```

```

        if(stringbuffer != null)
            vector.addElement(stringbuffer.toString());
        expr = new Object[vector.size()];
        vector.copyInto(expr);
    }

    static final char NOT[] = new char[2];
    static final Integer ANY = new Integer(0);
    static final Integer STAR = new Integer(1);
    Object expr[];
}

```

users.props file:

This example presents the format for the `users.props` file.

Note: The sample that is provided is intended to familiarize you with this feature. Do not use this sample in an actual production environment.

```

# 5639-D57, 5630-A36, 5630-A37, 5724-D18
# (C) COPYRIGHT International Business Machines Corp. 1997, 2005
# All Rights Reserved * Licensed Materials - Property of IBM
#
# Format:
# name:passwd:uid:gids:display name
# where name   = userId/userName of the user
#     passwd   = password of the user
#     uid      = uniqueId of the user
#     gid      = groupIds of the groups that the user belongs to
#     display name = a (optional) display name for the user.
bob:bob1:123:567:bob
dave:davel:234:678:
jay:jay1:345:678,789:Jay-Jay
ted:ted1:456:678:Teddy G
jeff:jeff1:222:789:Jeff
vikas:vikas1:333:789:vikas
bobby:bobby1:444:789:

```

groups.props file:

The following example illustrates the format for the `groups.props` file.

Note: The sample provided is intended to familiarize you with this feature. Do not use this sample in an actual production environment.

```

# 5639-D57, 5630-A36, 5630-A37, 5724-D18
# (C) COPYRIGHT International Business Machines Corp. 1997, 2005
# All Rights Reserved * Licensed Materials - Property of IBM
#
# Format:
# name:gid:users:display name
# where name   = groupId of the group
#     gid      = uniqueId of the group
#     users    = list of all the userIds that the group contains
#     display name = a (optional) display name for the group.
admins:567:bob:Administrative group
operators:678:jay,ted,dave:Operators group
users:789:jay,jeff,vikas,bobby:

```

Developing the UserRegistry interface for using custom registries

Implementing this interface enables WebSphere Application Server security to use custom registries. This capability extends the `java.rmi` file. With a remote registry, you can complete this process remotely.

About this task

Provide implementations of the following methods.

- Initialize the UserRegistry method, with initialize(java.util.Properties).

```
public void initialize(java.util.Properties props)
    throws CustomRegistryException,
           RemoteException;
```

This method is called to initialize the UserRegistry method. All the properties that are defined in the Custom User Registry panel propagate to this method.

For the FileRegistrySample.java sample file, the initialize method retrieves the names of the registry files that contain the user and group information.

This method is called during server bringup to initialize the registry. This method is also called when validation is performed by the administrative console, when security is on. This method remains the same as in Version 4.x.

- Authenticate users with checkPassword(String,String).

```
public String checkPassword(String userSecurityName, String password)
    throws PasswordCheckFailedException
           CustomRegistryException,
           RemoteException;
```

The checkPassword method is called to authenticate users when they log in using a name or user ID and a password. This method returns a string which, in most cases, is the user security name. A credential is created for the user for authorization purposes. This user name is also returned for the getCallerPrincipal enterprise bean call and the servlet calls the getUserPrincipal and getRemoteUser methods. See the getUserDisplayName method for more information if you have display names in your registry. In some situations, if you return a user other than the one who is logged in, you must verify that the user is valid in the registry.

For the FileRegistrySample.java sample file, the mapCertificate method gets the distinguished name (DN) from the certificate chain and makes sure it is a valid user in the registry before returning the user. For the sample, the checkPassword method checks the name and password combination in the user registry and, if they match, the method returns the user being authenticated.

This method is called for various scenarios, for example, by the administrative console to validate the user information after the user registry is initialized. This method is also called when you access protected resources in the product for authenticating the user and before proceeding with the authorization. This method is the same as in Version 4.x.

- Obtain user names from X.509 certificates with mapCertificate(X509Certificate[]).

```
public String mapCertificate(X509Certificate[] cert)
    throws CertificateMapNotSupportedException,
           CertificateMapFailedException,
           CustomRegistryException,
           RemoteException;
```

The mapCertificate method is called to obtain a user name from an X.509 certificate chain that is supplied by the browser. The complete certificate chain is passed to this method and the implementation can validate the chain if needed and get the user information. A credential is created for this user for authorization purposes. If browser certificates are not supported in your configuration, you can create the CertificateMapNotSupportedException exception. The consequence of not supporting certificates is authentication failure if the challenge type is certificates, even if valid certificates are in the browser.

This method is called when certificates are provided for authentication. For Web applications, when the authentication constraints are set to CLIENT-CERT in the web.xml file of the application, this method is called to map a certificate to a valid user in the registry. For Java clients, this method is called to map

the client certificates in the transport layer, when using transport layer authentication. When the identity assertion token, using the CSIV2 authentication protocol, is set to contain certificates, this method is called to map the certificates to a valid user.

In WebSphere Application Server Version 4.x, the input parameter is the X509Certificate certificate. In WebSphere Application Server Version 5.x and later, this parameter changes to accept an array of X509Certificate certificates such as a certificate chain. In Version 4.x, this parameter is called for Web applications only, but in version 5.x and later, you can call this method for both Web and Java clients.

- Obtain the security realm name with `getRealm()`.

```
public String getRealm()
    throws CustomRegistryException,
           RemoteException;
```

The `getRealm` method is called to get the name of the security realm. The name of the realm identifies the security domain for which the registry authenticates users. If this method returns a null value, a `customRealm` default name is used.

For the `FileRegistrySample.java` sample file, the `getRealm` method returns the `customRealm` string. One of the calls to this method occurs when the user registry information is validated. This method is the same method as in Version 4.x.

- Obtain the list of users from the registry with `getUsers(String,int)`.

```
public Result getUsers(String pattern, int limit)
    throws CustomRegistryException,
           RemoteException;
```

The `getUsers` method returns the list of users from the registry. The names of users depend on the pattern parameter. The number of users are limited by the limit parameter. In a registry that has many users, getting all the users is not practical. So the limit parameter is introduced to limit the number of users retrieved from the registry. A limit of zero (0) indicates to return all the users that match the pattern and might cause problems for large registries. Use this limit with care.

The custom registry implementations are expected to support at least the wildcard search (*). For example, a pattern of asterisk (*) returns all the users and a pattern of (b*) returns the users starting with *b*.

The return parameter is an object with a `com.ibm.websphere.security.Result` type. This object contains two attributes, a `java.util.List` and a `java.lang.Boolean` attribute. The list contains the users that are returned and the Boolean flag indicates if more users are available in the user registry for the search pattern. This Boolean flag is used to indicate to the client whether more users are available in the registry.

In the `FileRegistrySample.java` sample file, the `getUsers` method retrieves the required number of users from the user registry and sets them as a list in the Result object. To find out if more users are presented than requested, the sample gets one more user than requested and if it finds the additional user, it sets the Boolean flag to `true`. For pattern matching, the `match` method in the `RegExpSample` class is used, which supports wildcard characters such as the asterisk (*) and the question mark (?).

This method is called by the administrative console to add users to roles in the various map-users-to-roles panels. The administrative console uses the Boolean set in the Result object to indicate that more entries matching the pattern are available in the user registry.

In WebSphere Application Server Version 4.x, this method specifies to take only the pattern parameter. The return is a list. In WebSphere Application Server Version 5.x or later, this method is changed to take one additional parameter, the limit. Ideally, your implementation changes to take the limit value and limits the users that are returned. The return is changed to return a Result object, which consists of the list and a flag that indicates if more entries exist. When the list returns, use the `Result.setList(List)` method to set the list in the Result object. If more entries exist than requested in the limit parameter, set the Boolean attribute to `true` in the result object, using the `Result.setHasMore` method. The default for the Boolean attribute in the result object is `false`.

- Obtain the display name of a user with `getUserDisplayName(String)`.

```
public String getUserDisplayName(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

The `getUserDisplayName` method returns a display name for a user, if one exists. The display name is an optional string that describes the user that you can set in some registries. This descriptive name is for the user and does not have to be unique in the registry.

For example in Windows systems, you can display the full name of the user.

If you do not need display names in your registry, return null or an empty string for this method.

If display names existed for any user in WebSphere Application Server Version 4.x, these names were useful for the Enterprise JavaBean (EJB) method call `getCallerPrincipal` and the servlet calls `getUserPrincipal` and `getRemoteUser`. If the display names are not the same as the security name for any user, the display names are returned for the previously mentioned enterprise beans and servlet methods. Returning display names for these methods might become problematic in some situations because the display names might not be unique in the user registry. Avoid this problem by changing the default behavior to return the user security name instead of the user display name in this version of the product. For more information on how to set properties for the custom registry, see the section on *Setting Properties for Custom Registries*.

In the `FileRegistrySample.java` sample file, this method returns the display name of the user whose name matches the user name that is provided. If the display name does not exist, this method returns an empty string.

This method can be called by the product to present the display names in the administrative console, or by using the command line and the **wsadmin** tool. Use this method for display purposes only. This method is the same as in Version 4.x.

- Obtain the unique ID of a user with `getUniqueUserId(String)`.

```
public String getUniqueUserId(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns the unique ID of the user, given the security name.

In the `FileRegistrySample.java` sample file, this method returns the `uniqueUserId` value of the user whose name matches the supplied name. This method is called when forming a credential for a user and also when creating the authorization table for the application.

- Obtain the security name of a user with `getUserSecurityName(String)`.

```
public String getUserSecurityName(String uniqueUserId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns the security name of a user given the unique ID. In the `FileRegistrySample.java` sample file, this method returns the security name of the user whose unique ID matches the supplied ID.

This method is called to make sure a valid user exists for a given `uniqueUserId`. This method is called to get the security name of the user when the `uniqueUserId` is obtained from a token.

- Check whether a given user is a valid user in the registry with `isValidUser(String)`.

```
public boolean isValidUser(String userSecurityName)
    throws CustomRegistryException,
           RemoteException;
```

This method indicates whether the given user is a valid user in the registry.

In the `FileRegistrySample.java` sample file, this method returns `true` if the user is found in the registry, otherwise this method returns `false`. This method is primarily called in situations where knowing if the

user exists in the directory prevents problems later. For example, in the `mapCertificate` call, when the name is obtained from the certificate if the user is not found as a valid user in the user registry, you can avoid trying to create the credential for the user.

- Return the list of groups from the user registry with `getGroups(String,int)`.

```
public Result getGroups(String pattern, int limit)
    throws CustomRegistryException,
           RemoteException;
```

The `getGroups` method returns the list of groups from the user registry. The names of groups depend on the pattern parameter. The number of groups is limited by the limit parameter. In a registry that has many groups, getting all the groups is not practical. So, the limit parameter is introduced to limit the number of groups retrieved from the user registry. A limit of zero (0) implies to return all the groups that match the pattern and can cause problems for large user registries. Use this limit with care. The custom registry implementations are expected to support at least the wildcard search (*). For example, a pattern of asterisk (*) returns all the users and a pattern of (b*) returns the users starting with *b*.

The return parameter is an object of the `com.ibm.websphere.security.Result` type. This object contains the `java.util.List` and `java.lang.Boolean` attributes. The list contains the groups that are returned and the Boolean flag indicates whether more groups are available in the user registry for the pattern searched. This Boolean flag is used to indicate to the client if more groups are available in the registry.

In the `FileRegistrySample.java` sample file, the `getUsers` method retrieves the required number of groups from the user registry and sets them as a list in the Result object. To find out if more groups are presented than requested, the sample gets one more user than requested and if it finds the additional user, it sets the Boolean flag to `true`. For pattern matching, the `match` method in the `RegExpSample` class is used, which supports the asterisk (*) and question mark (?) characters.

This method is called by the administrative console to add groups to roles in the various `map-groups-to-roles` panels. The administrative console uses the boolean set in the Result object to indicate that more entries matching the pattern are available in the user registry.

In WebSphere Application Server Version 4, this method is used to take the pattern parameter only and returns a list. In WebSphere Application Server Version 5.x or later, this method is changed to take the limit parameter. Change to take the limit value and limit the users that are returned. The return is changed to return a Result object, which consists of the list and a flag that indicates whether more entries exist. Use the `Result.setList(List)` method to set the list in the Result object. If more entries exist than requested in the limit parameter, set the Boolean attribute to `true` in the Result object using the `Result.setHasMore` method. The default for the Boolean attribute in the Result object is `false`.

- Obtain the display name of a group with `getGroupDisplayName(String)`.

```
public String getGroupDisplayName(String groupSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

The `getGroupDisplayName` method returns a display name for a group if one exists. The display name is an optional string that describes the group that you can set in some user registries. This name is a descriptive name for the group and does not have to be unique in the registry. If you do not need to have display names for groups in your registry, return null or an empty string for this method.

In the `FileRegistrySample.java` sample file, this method returns the display name of the group whose name matches the group name that is provided. If the display name does not exist, this method returns an empty string.

The product can call this method to present the display names in the administrative console or through the command line using the **wsadmin** tool. This method is used for display purposes only.

- Obtain the unique ID of a group with `getUniqueGroupId(String)`.

```
public String getUniqueGroupId(String groupSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns the unique ID of the group that is given the security name.

In the `FileRegistrySample.java` sample file, this method returns the security name of the group whose unique ID matches the supplied ID. This method verifies that a valid group exists for a given `uniqueGroupId` ID.

- Obtain the unique IDs of all groups to which a user belongs with `getUniqueGroupIds(String)`.

```
public List getUniqueGroupIds(String uniqueUserId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns the unique IDs of all the groups to which a user belongs.

In the `FileRegistrySample.java` sample file, this method returns the unique ID of all the groups that contain this `uniqueUserId` ID. This method is called when creating the credential for the user. As part of creating the credential, all the `groupUniqueIds` IDs in which the user belongs are collected and put in the credential for authorization purposes when groups are given access to a resource.

- Obtain the security name of a group with `getGroupSecurityName(String)`.

```
public String getGroupSecurityName(String uniqueGroupId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns the security name of a group given its unique ID.

In the `FileRegistrySample.java` sample file, this method returns the security name of the group whose unique ID matches the supplied ID. This method verifies that a valid group exists for a given `uniqueGroupId` ID.

- Determine whether a group is a valid group in the registry with `isValidGroup(String)`.

```
public boolean isValidGroup(String groupSecurityName)
    throws CustomRegistryException,
           RemoteException;
```

This method indicates if the given group is a valid group in the registry.

In the `FileRegistrySample.java` sample file, this method returns `true` if the group is found in the registry, otherwise the method returns `false`. This method can be used in situations where knowing whether the group exists in the directory might prevent problems later.

- Obtain all groups to which a user belongs with `getGroupsForUser(String)`.

```
public List getGroupsForUser(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;
```

This method returns all the groups to which a user belongs whose name matches the supplied name.

This method is similar to the `getUniqueGroupIds` method with the exception that the security names are used instead of the unique IDs.

In the `FileRegistrySample.java` sample file, this method returns all the group security names that contain the `userSecurityName` name.

This method is called by the administrative console or the scripting tool to verify that the users entered for the `RunAs` roles are already part of that role in the users and groups-to-role mapping. This check is required to ensure that a user cannot be added to a `RunAs` role unless that user is assigned to the role in the users and groups-to-role mapping either directly or indirectly through a group that contains this user. Because a group in which the user belongs can be part of the role in the users and groups-to-role mapping, this method is called to check if any of the groups that this user belongs to mapped to that role.

- Retrieve users from a specified group with `getUsersForGroup(String,int)`.

```

public Result getUsersForGroup(String groupSecurityName, int limit)
    throws NotImplementedException,
        EntryNotFoundException,
        CustomRegistryException,
        RemoteException;

```

This method retrieves users from the specified group. The number of users returned is limited by the limit parameter. A limit of zero (0) indicates to return all of the users in that group. This method is not directly called by the WebSphere Application Server security component. However, this method can be called by other components. In rare situations, if you are working with a user registry where getting all the users from any of your groups is not practical, you can create the `NotImplementedException` exception for the particular groups. In this case, verify that if the process choreographer is installed the staff assignments are not modeled using these particular groups. If no concern exists about returning the users from groups in the user registry, it is recommended that you do not create the `NotImplemented` exception when implementing this method.

The return parameter is an object with a `com.ibm.websphere.security.Result` type. This object contains the `java.util.List` and `java.lang.Boolean` attributes. The list contains the users that are returned and the Boolean flag, which indicates whether more users are available in the user registry for the search pattern. This Boolean flag indicates to the client whether users are available in the user registry.

In the example, this method gets one user more than the requested number of users for a group, if the limit parameter is not set to zero (0). If the method succeeds in getting one more user, the Boolean flag is set to `true`.

In WebSphere Application Server Version 4, this `getUsers` method is mandatory for the product. For WebSphere Application Server Version 5.x or later, this method can create the `NotImplementedException` exception in situations where it is not practical to get the requested set of users. However, create this exception in rare situations when as other components can be affected. In Version 4, this method accepts only the pattern parameter and returns a list. In Version 5, this method accepts the limit parameter. Change your implementation to take the limit value and limit the users that are returned. The return changes to return a `Result` object, which consists of the list and a flag that indicates whether more entries exist. When the list is returned, use the `Result.setList(List)` method to set the list in the `Result` object. If more entries than requested are in the limit parameter, set the Boolean attribute to `true` in the `Result` object using `Result.setHasMore` method. The default for the Boolean attribute in the `Result` object is `false`.

- Implement the `createCredential(String)` method.

Note: The first two lines of the following code sample are split for illustrative purposes only.

```

public com.ibm.websphere.security.cred.WSCredential createCredential(String userSecurityName)
    throws NotImplementedException,
        EntryNotFoundException,
        CustomRegistryException,
        RemoteException;

```

In this release the WebSphere Application Server, the `createCredential` method is not called. You can return `null`. In the example, a `null` value is returned.

What to do next

Managing the realm in a federated repository configuration

Follow this topic to manage the realm in a federated repository configuration.

Before you begin

The realm can consist of identities in:

- The file-based repository that is built into the system
- One or more external repositories

- Both the built-in, file-based repository and in one or more external repositories

Before you configure your realm, review “Federated repositories limitations” on page 173.

1. Configure your realm by using one of the following topics. You might be configuring your realm for the first time or changing an existing realm configuration.
 - “Using a single built-in, file-based repository in a new configuration under Federated repositories” on page 176
 - “Changing a federated repository configuration to include a single built-in, file-based repository only” on page 181
 - “Configuring a single, Lightweight Directory Access Protocol repository in a new configuration under Federated repositories” on page 182
 - “Changing a federated repository configuration to include a single, Lightweight Directory Access Protocol repository only” on page 183
 - “Configuring multiple Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 184
 - “Configuring a single built-in, file-based repository and one or more Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 185
2. Configure supported entity types using the steps described in “Configuring supported entity types in a federated repository configuration” on page 213. You must configure supported entity types before you can manage this account with Users and Groups. The Base entry for the default parent determines the repository location where entities of the specified type are placed on a create operation.
3. Optional: Use one or more of the following tasks to extend the capabilities of storing data and attributes in your realm:
 - a. Configure an entry mapping repository using the steps described in “Configuring an entry mapping repository in a federated repository configuration” on page 210. An entry mapping repository is used to store data for managing profiles on multiple repositories.
 - b. Configure a property extension repository using the steps described in “Configuring a property extension repository in a federated repository configuration” on page 197. A property extension repository is used to store attributes that cannot be stored in your Lightweight Directory Access Protocol (LDAP) server.
 - a. Set up a database repository using wsadmin commands as described in “Setting up an entry mapping repository, a property extension repository, or a custom registry database repository using wsadmin commands” on page 202
4. Optional: Use one or more of the following advanced user tasks to extend the capabilities of LDAP repositories in your realm:
 - “Increasing the performance of the federated repository configuration” on page 218
 - “Configuring Lightweight Directory Access Protocol entity types in a federated repository configuration” on page 228
 - “Configuring group attribute definition settings in a federated repository configuration” on page 230
5. Optional: Manage repositories that are configured in your system by following the steps described in “Managing repositories in a federated repository configuration” on page 215.
6. Optional: Add an external repository into your realm by following the steps described in “Adding an external repository in a federated repository configuration” on page 196.
7. Optional: Change the password for the repository that is configured under federated repositories by the following steps described in “Changing the password for a repository under a federated repositories configuration” on page 174.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm

definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.

2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Realm configuration settings

Use this page to manage the realm. The realm can consist of identities in the file-based repository that is built into the system, in one or more external repositories, or in both the built-in, file-based repository and one or more external repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

A single built-in, file-based repository is built into the system and included in the realm by default.

You can configure one or more Lightweight Directory Access Protocol (LDAP) repositories to store identities in the realm. Click **Add base entry to realm** to specify a repository configuration and a base entry into the realm. You can configure multiple different base entries into the same repository.

Click **Remove** to remove selected repositories from the realm. Repository configurations and contents are not destroyed. The following restrictions apply:

- The realm must always contain at least one base entry; therefore, you cannot remove every entry.
- If you plan to remove the built-in, file-based repository from the administrative realm, verify that at least one user in another member repository is a console user with administrative rights. Otherwise, you must disable security to regain access to the administrative console.

WebSphere Application Server Version 6.1 distinguishes between the user identities for administrators who manage the environment and server identities for authenticating server to server communications. In most cases, server identities are automatically generated and are not stored in a repository.

However, if you are adding a Version 5.0.x or 6.0.x node to a Version 6.1 cell, you must ensure that the Version 5.x or Version 6.0.x server identity and password are defined in the repository for this cell. Enter the server user identity and password on this panel.

Realm name:

Specifies the name of the realm. You can change the realm name.

Primary administrative user name:

Specifies the name of the user with administrative privileges that is defined in the repository, for example, adminUser.

The user name is used to log on to the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Automatically generated server identity:

Enables the application server to generate the server identity, which is recommended for environments that contain only Version 6.1 or later nodes. Automatically generated server identities are not stored in a user repository.

You can change this server identity on the Authentication mechanisms and expiration panel. To access the Authentication mechanisms and expiration panel, click **Security > Global security > Authentication mechanisms and expiration**. Change the value of the Internal server ID field.

Default: Enabled

Server identity that is stored in the repository:

Specifies a user identity in the repository that is used for internal process communication. Cells that contain Version 5.x or 6.0.x nodes require a server user identity that is defined in the active user repository.

Default: Enabled

Server user ID or administrative user on a Version 6.0.x node:

Specifies the user ID that is used to run the application server for security purposes.

Password:

Specifies the password that corresponds to the server ID.

Ignore case for authorization:

Specifies that a case-insensitive authorization check is performed.

If case sensitivity is not a consideration for authorization, enable the **Ignore case for authorization** option.

Base entry:

Specifies the base entry within the realm. This entry and its descendents are part of the realm.

Repository identifier:

Specifies a unique identifier for the repository. This identifier uniquely identifies the repository within the cell.

Repository type:

Specifies the repository type, such as File or LDAP.

Federated repositories limitations

This topic outlines known limitations and important information for configuring federated repositories.

Configuring federated repositories in a mixed-version environment

In a mixed-version deployment manager cell that contains both Version 6.1.x and Version 5.x or 6.0.x nodes, the following limitations apply for configuring federated repositories:

- You can configure only one Lightweight Directory Access Protocol (LDAP) repository under federated repositories, and the repository must be supported by Version 5.x or 6.0.x.
- You can specify a realm name that is compatible with prior versions only. The host name and the port number represent the realm for the LDAP server in a mixed-version nodes cell. For example, machine1.austin.ibm.com:389.
- You must configure a stand-alone LDAP registry; the LDAP information in both the stand-alone LDAP registry and the LDAP repository under the federated repositories configuration must match. During node synchronization, the LDAP information from the stand-alone LDAP registry propagates to the Version 5.x or 6.0.x nodes.

Note: Before node synchronization, verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. Do not set the stand-alone LDAP registry as the current realm definition.

- You cannot configure an entry mapping repository or a property extension repository in a mixed-version deployment manager cell.

Configuring LDAP servers in a federated repository

The LDAP connection connectTimeout default value is 20 seconds. LDAP should respond within 20 seconds for any request from WebSphere Application Server. If you cannot connect to your LDAP within this time, make sure that your LDAP is running. A connection error displays at the top of the LDAP configuration panel when the connection timeout exceeds 20 seconds.

Coexisting with Tivoli Access Manager

For Tivoli Access Manager to coexist with a federated repositories configuration, the following limitations apply:

- You can configure only one LDAP repository under federated repositories, and that LDAP repository configuration must match the LDAP server configuration under Tivoli Access Manager.
- The distinguished name for the realm base entry must match the LDAP distinguished name (DN) of the base entry within the repository. In WebSphere Application Server, Tivoli Access Manager recognizes the LDAP user ID and LDAP DN for both authentication and authorization. The federated repositories configuration does not include additional mappings for the LDAP user ID and DN.
- The federated repositories functionality does not recognize the metadata that is specified by Tivoli Access Manager. When users and groups are created under user and group management, they are not formatted using the Tivoli Access Manager metadata. The users and groups must be manually imported into Tivoli Access Manager before you use them for authentication and authorization.

Changing the password for a repository under a federated repositories configuration

Passwords allow security control over the repositories under a federated repositories configuration. As part of managing the realm in a federated repository configuration, one of the optional tasks you can perform is to change the password of an individual repository that is under a federated repositories configuration.

Before you begin

Before you change the password for the repository that is configured under federated repositories, ensure that the WebSphere Application Server is running and the target repository for the password change is configured under the federated repositories configuration. In a Network Deployment environment, you must complete these steps on the deployment manager.

- Changing the password for a repository using the dynamic **updateIdMgrLDAPBindInfo** command Use the following steps to change the Lightweight Directory Access Protocol (LDAP) bind distinguished name (DN) or bind password of an LDAP repository.
From a **wsadmin** prompt, you can enter the following command to display a list of arguments for the **updateIdMgrLDAPBindInfo** command: `$AdminTask help UpdateIdMgrLDAPBindInfo`
 1. Start the wsadmin command-line utility. The wsadmin command is found in the `app_server_root/bin` directory. The WebSphere Application Server and wsadmin must remaining running.
 2. Use an LDAP tool to change the password of the LDAP repository. Some LDAP repositories require a stop and start of the LDAP server to change the password.
 3. From the **wsadmin** prompt, enter the **updateIdMgrLDAPBindInfo** command to update the LDAP password under the federated repository. The change is also reflected in the `wimconfig.xml` file.
- Changing the password for a repository using the **updateIdMgrDBRepository** command
 1. Start the wsadmin command-line utility. The wsadmin command is found in the `app_server_root/bin` directory. The wsadmin command session must remain running. If WebSphere Application Server is not started, you need to open a wsadmin command session in local mode. `wsadmin -conntype none`
 2. Log in to the Administrative Console for WebSphere Application Server.
 3. Change the password for the repository.
 4. From the Administrative Console, change the data source (J2C) password. You access the proper console page by clicking **Resources > JDBC > Data sources > data_source > JAAS - J2C authentication data**.
 5. From the Administrative Console, save your changes to the master configuration.
 6. From the **wsadmin** prompt, use the **updateIdMgrDBRepository** command to update the password in the `wimconfig.xml` file.
 7. From the **wsadmin** prompt, save your changes to the master configuration. The following command is used to save the master configuration: `$AdminConfig save`.
 8. Restart the WebSphere Application Server.
- Changing the password for a repository using the **setIdMgrPropertyExtensionRepository** command
 1. Start the wsadmin command-line utility. The wsadmin command is found in the `app_server_root/bin` directory. The wsadmin command session must remain running. If WebSphere Application Server is not started, you need to open a wsadmin command session in local mode.
`wsadmin -conntype none`
 2. Log in to the Administrative Console for WebSphere Application Server.
 3. Change the password for the repository.
 4. From the Administrative Console, change the data source (J2C) password. You access the proper console page by clicking **Resources > JDBC > Data sources > data_source > JAAS - J2C authentication data**.
 5. From the Administrative Console, save your changes to the master configuration.
 6. From the **wsadmin** prompt, use the **setIdMgrPropertyExtensionRepository** command to update the password in the `wimconfig.xml` file.
 7. From the **wsadmin** prompt, save your changes to the master configuration. The following command is used to save the master configuration: `$AdminConfig save`.
 8. Restart the WebSphere Application Server.
- Changing the password for a repository using the **setIdMgrEntryMappingRepository** command
 1. Start the wsadmin command-line utility. The wsadmin command is found in the `app_server_root/bin` directory. The wsadmin command session must remain running. If WebSphere Application Server is not started, you need to open a wsadmin command session in local mode.
`wsadmin -conntype none`
 2. Log in to the Administrative Console for WebSphere Application Server.
 3. Change the password for the repository.

4. From the Administrative Console, change the data source (J2C) password. You access the proper console page by clicking **Resources > JDBC > Data sources > data_source > JAAS - J2C authentication data**.
 5. From the Administrative Console, save your changes to the master configuration.
 6. From the **wsadmin** prompt, use the **setIdMgrEntryMappingRepository** command to update the password in the wimconfig.xml file.
 7. From the **wsadmin** prompt, save your changes to the master configuration. The following command is used to save the master configuration: `$AdminConfig save`.
 8. Restart the WebSphere Application Server.
- Changing the password for a repository using the **updateIdMgrLDAPServer** command
 1. Start the wsadmin command-line utility. The wsadmin command is found in the `app_server_root/bin` directory. The wsadmin command session must remain running. If WebSphere Application Server is not started, you need to open a wsadmin command session in local mode.


```
wsadmin -conntype none
```
 2. Change the password for the repository.
 3. From the **wsadmin** prompt, use the **updateIdMgrDBRepository** command to update the password in the wimconfig.xml file.
 4. From the **wsadmin** prompt, save your changes to the master configuration. The following command is used to save the master configuration: `$AdminConfig save`.
 5. Restart the WebSphere Application Server.

Results

The password for the repository has been changed.

Using a single built-in, file-based repository in a new configuration under Federated repositories

Follow this task to use a single built-in, file-based repository in a new configuration under Federated repositories.

Before you begin

To use the default configuration under Federated repositories that includes a single built-in, file-based repository only, you need to know the primary administrative user name of the user who manages WebSphere Application Server resources and user accounts.

Note:

Client certificate login is not supported in a realm that includes a single built-in, file-based repository or a single built-in, file-based repository with other repositories.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Leave the Realm name field value as `defaultWIMFileBasedRealm`.
4. Enter the name of the primary administrative user in the Primary administrative user name field, for example, `adminUser`.
5. Leave the **Ignore case for authorization** option enabled.
6. Click **OK**.
7. Provide an administrative user password. This panel displays only when a built-in, file-based repository is included in the realm. Otherwise, the panel does not display. If a built-in, file-based repository is included, complete the following steps:

- a. Supply a password for the administrative user in the Password field.
- b. Confirm the password of the primary administrative user in the Confirm password field.
- c. Click **OK**.

Results

After completing these steps, your new configuration under Federated repositories includes a single built-in, file-based repository only.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps, as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Administrative user password settings:

Use this page to set a password for the administrative user who manages the product resources and user accounts.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. If your federated repository configuration includes a built-in, file-based repository, then the **Administrative user password** panel displays when changes are applied.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Password:

Specifies the password of the administrative user who manages the product resources and user accounts.

Confirm password:

Confirms the password of the administrative user who manages the product resources and user accounts.

Federated repository wizard settings:

Use this security wizard page to complete the basic requirements to connect the application server to a federated repository.

To view this security wizard page, complete the following steps

1. Click **Security > Global security > Security configuration wizard**.
2. Select your protection settings and click **Next**.
3. Select the **Federated repositories** option and click **Next**.

You can modify your federated repository configuration by completing the following steps:

1. Click **Security > Global security**.
2. Under User account repository, select Federated repository and click **Configure**.

Note: This wizard is used for the initial configuration of a built-in, file-based repository. The user name and password do not have to be in the federated repository because they will be created. If you have previously configured federated repositories, do not use the Security configuration wizard to modify your configuration. Instead, modify your configuration using the Federated repositories selection under User account repository on the Global security panel.

Primary administrative user name:

Specifies the name of the user with administrative privileges that is defined in the repository, for example, adminUser.

The user name is used to log on to the administrative console when administrative security is enabled. Version 6.1 requires an administrative user that is distinct from the server user identity so that administrative actions can be audited.

Note: In WebSphere Application Server, Versions 5.x and 6.0.x, a single user identity is required for both administrative access and internal process communication. When migrating to Version 6.1, this identity is used as the server user identity. You need to specify another user for the administrative user identity.

Password:

Specifies the password of the administrative user who manages the product resources and user accounts.

Confirm password:

Confirms the password of the administrative user who manages the product resources and user accounts.

Configuring a single built-in, file-based repository in a new configuration under federated repositories using wsadmin

You can use the Jython or Jacl scripting language with the wsadmin tool to configure a single built-in, file-based repository in a new configuration under Federated repositories.

Before you begin

Shut down the application server and ensure you have the primary administrator id and password.

About this task

The federated repositories configuration file, wimconfig.xml, is supported by WebSphere Application Server 6.1.x and is located in the *app_server_root/profiles/profile_name/config/cells/cell_name/wim/config* directory.

Use the following steps to configure for use a single built-in, file-based repository in a new configuration for federated repositories.

1. Launch the wsadmin scripting tool using the Jython scripting language.

2. Create the fileRegistry.xml file, which is the user registry itself, if it does not already exist. If the fileRegistry.xml file does exist, this step just adds the user to registry.

Using Jython:

```
AdminTask.addFileRegistryAccount('-userId isoet01s01 -password oets01')
```

Using Jacl:

```
$AdminTask addFileRegistryAccount {-userId isoet01s01 -password oets01}
```

3. Update the security.xml file to enable administrative security, set the activeUserRegistry to use federated repositories, and update the primaryAdmin and its password.

Using Jython:

```
AdminTask.applyWizardSettings('-secureApps false  
-secureLocalResources false  
-userRegistryType WIMUserRegistry  
-customRegistryClass com.ibm.ws.wim.registry.WIMUserRegistry  
-adminName isoet01s01 -adminPassword oets01')
```

Using Jacl:

```
$AdminTask applyWizardSettings {-secureApps false  
-secureLocalResources false  
-userRegistryType WIMUserRegistry  
-customRegistryClass com.ibm.ws.wim.registry.WIMUserRegistry  
-adminName isoet01s01  
-adminPassword oets01}
```

4. Save your configuration changes. Enter the following commands to save the new configuration and close the wsadmin tool:

Using Jython:

```
AdminConfig.save()
```

Using Jacl:

```
$AdminConfig save
```

5. Restart the application server.

FileRegistryCommands command group for the AdminTask object:

Federated repositories provides a file registry. Use the commands in the FileRegistryCommands command group to administer the file registry using the wsadmin tool.

Use the following commands in the FileRegistryCommands group to modify the federated repository file registry:

- “addFileRegistryAccount command”
- “changeFileRegistryAccountPassword command” on page 180

addFileRegistryAccount command

The addFileRegistryAccount command adds an account to the file registry. You must save your configuration changes after running this command to save the new account to the master repository.

Target object

None

Required parameters

-userId

Specifies the ID of the user to add to the file registry. (String, required)

-password

Specifies the password of the user to add to the file registry. (String, required)

Optional parameters

-parent

Specifies the parent of the entity. (String, optional)

.

Return value

This command returns a message that indicates that the command ran successfully, as the following example displays:

```
'CWWIM4544I Account newAcct(uid=newAcct,o=defaultWIMFileBasedRealm) is stored in the file registry in the temporary workspace. You must use the "$AdminConfig save" command to save it in the master repository.'
```

Batch mode example usage

- Using Jython string:

```
AdminTask.addFileRegistryAccount('-userId newAcct -password new22password')
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.addFileRegistryAccount(['-userId', 'newAcct', '-password', 'new22password'])
```

changeFileRegistryAccountPassword command

The `changeFileRegistryAccountPassword` changes the password for the file registry account.

Target object

None.

Required parameters

-userId

Specifies the user ID of interest. (String, required)

-password

Specifies the new password. (String, required)

Optional parameters

-uniqueName

Specifies the fully-qualified unique name of the administrator. (String, optional)

Return value

This command returns a message that indicates that the command ran successfully, as the following example displays:

```
'CWWIM4545I The password is changed for newAcct(uid=newAcct,o=defaultWIMFileBasedRealm) in the file registry in the temporary workspace. You must use the "$AdminConfig save" command to save it in the master repository.'
```

Batch mode example usage

- Using Jython string:

```
AdminTask.changeFileRegistryAccountPassword('-userId newAcct -password newPassword -uniqueName uid=newAcct,o=defaultWIMFileBasedRealm')
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.changeFileRegistryAccountPassword(['-userId', 'newAcct', '-password', 'newPassword', '-uniqueName', 'uid=newAcct,o=defaultWIMFileBasedRealm'])
```

Changing a federated repository configuration to include a single built-in, file-based repository only

Follow this task to change your federated repository configuration to include a single built-in, file-based repository only.

Before you begin

To change your federated repository configuration to include a single built-in, file-based repository only, you need to know the primary administrative user name of the user who manages WebSphere Application Server resources and user accounts.

Note: Client certificate login is not supported in a realm that includes a single built-in, file-based repository or a single built-in, file-based repository with other repositories.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Enter the name of the realm in the Realm name field. If the realm contains a single built-in, file-based repository only, you must specify `defaultWIMFileBasedRealm` as the realm name.
4. Enter the name of the primary administrative user in the Primary administrative user name field, for example, `adminUser`.
5. Enable the **Ignore case for authorization** option.
6. Click **Use built-in repository** if the built-in, file-based repository is not listed in the collection.
7. Select all repositories in the collection that are not of type File and click **Remove**.
8. Click **OK**.
9. Provide an administrative user password. This panel displays only when a built-in, file-based repository is included in the realm. Otherwise, it does not display. If a built-in, file-based repository is included, complete the following steps:
 - a. Supply a password for the primary administrative user in the Password field.
 - b. Confirm the password of the primary administrative user in the Confirm password field.
 - c. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes a single built-in, file-based repository only, is configured.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.

3. If you are enabling security, complete the remaining steps, as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Configuring a single, Lightweight Directory Access Protocol repository in a new configuration under Federated repositories

Follow this task to configure a single, Lightweight Directory Access Protocol (LDAP) repository in a new configuration under Federated repositories.

Before you begin

To configure an LDAP repository in a new configuration under Federated repositories, you must know a valid user name (ID), the user password, the server host and port and, if necessary, the bind distinguished name (DN) and the bind password. You can choose any valid user in the repository that is searchable. In some LDAP servers, administrative users are not searchable and cannot be used (for example, cn=root in SecureWay). This user is referred to as the WebSphere Application Server *administrative user name* or *administrative ID* in the documentation. Being an administrative ID means a user has special privileges when calling some protected internal methods. Normally, this ID and password are used to log in to the administrative console after you turn on security. You can use other users to log in, if those users are part of the administrative roles.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. On the Federated repositories panel, complete the following steps:
 - a. Enter the name of the realm in the Realm name field. You can change the existing realm name.
 - b. Enter the name of the primary administrative user in the Primary administrative user name field, for example, adminUser.
 - c. Optional: Select the **Ignore case for authorization** option. When you enable this option, the authorization check is case-insensitive. Normally, an authorization check involves checking the complete DN of a user, which is unique in the realm and is case-insensitive. Clear this option when all of the member repositories in the realm are case-sensitive.

Note: Some repositories contain data that is case-sensitive only, and some repositories contain data that is case-insensitive only. Do not include both case-sensitive and case-insensitive repositories in the realm. For example, do not include case-sensitive repositories in the realm with a built-in, file-based repository.

- d. Click **Add base entry to realm** to add a base entry that uniquely identifies the external repository in the realm. Then complete the steps in “Adding an external repository in a federated repository configuration” on page 196.
4. On the Federated repositories panel, complete the following steps:
 - a. Select the built-in, file-based repository in the collection, and click **Remove**.

Note: Before you remove the built-in, file-based repository from the administrative realm, verify that at least one user in another member repository is a console user with administrative rights. Otherwise, you must disable security to regain access to the administrative console.

- b. Click **OK**.

Results

After completing these steps, your new configuration under Federated repositories includes a single, LDAP repository only.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Changing a federated repository configuration to include a single, Lightweight Directory Access Protocol repository only

Follow this task to change your federated repository configuration to include a single, Lightweight Directory Access Protocol repository (LDAP) repository only.

Before you begin

To configure an LDAP repository in a federated repository configuration, you must know a valid user name (ID), the user password, the server host and port and, if necessary, the bind distinguished name (DN) and the bind password. You can choose any valid user in the repository that is searchable. In some LDAP servers, administrative users are not searchable and cannot be used (for example, cn=root in SecureWay). This user is referred to as a WebSphere Application Server *administrative user name* or *administrative ID* in the documentation. Being an administrative ID means a user has special privileges when calling some protected internal methods. Normally, this ID and password are used to log into the administrative console after you turn on security. You can use other users to log in if those users are part of the administrative roles.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Enter the name of the realm in the Realm name field. You can change the existing realm name.
4. Enter the name of the primary administrative user in the Primary administrative user name field, for example, adminUser.
5. Optional: Select the **Ignore case for authorization** option. When you enable this option, the authorization check is case-insensitive. Normally, an authorization check involves checking the complete DN of a user, which is unique in the realm and is case-insensitive. Clear this option when all of the member repositories in the realm are case-sensitive.

Note: Some repositories contain data that is case-sensitive only, and some repositories contain data that is case-insensitive only. Do not include both case-sensitive and case-insensitive repositories in the realm. For example, do not include case-sensitive repositories in the realm with a built-in, file-based repository.

6. Optional: Click **Add base entry to realm** if the LDAP repository that you need is not contained in the collection. Then complete the steps in “Adding an external repository in a federated repository configuration” on page 196.
7. On the Federated repositories panel, complete the following steps:
 - a. Optional: Select the repositories in the collection that you do not need in the realm and click **Remove**.

Note: The realm must always contain at least one base entry; therefore, you cannot remove every entry.

- b. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes a single LDAP repository only, is configured.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Configuring multiple Lightweight Directory Access Protocol repositories in a federated repository configuration

Follow this task to configure multiple Lightweight Directory Access Protocol (LDAP) repositories in a federated repository configuration.

Before you begin

To configure an LDAP repository in a federated repository configuration, you must know a valid user name (ID), the user password, the server host and port and, if necessary, the bind distinguished name (DN) and the bind password. You can choose any valid user in the repository that is searchable. In some LDAP servers, administrative users are not searchable and cannot be used (for example, cn=root in SecureWay). This user is referred to as a WebSphere Application Server *administrative user name* or *administrative ID* in the documentation. Being an administrative ID means a user has special privileges when calling some protected internal methods. Normally, this ID and password are used to log into the administrative console after you turn on security. You can use other users to log in if those users are part of the administrative roles.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Enter the name of the realm in the Realm name field. You can change the existing realm name.
4. Enter the name of the primary administrative user in the Primary administrative user name field, for example, adminUser.
5. Optional: Select the **Ignore case for authorization** option. When you enable this option, the authorization check is case-insensitive. Normally, an authorization check involves checking the complete DN of a user, which is unique in the realm and is case-insensitive. Clear this option when all of the member repositories in the realm are case-sensitive.

Note: Some repositories contain data that is case-sensitive only, and some repositories contain data that is case-insensitive only. Do not include both case-sensitive and case-insensitive repositories in the realm. For example, do not include case-sensitive repositories in the realm with a built-in, file-based repository.

6. Optional: Click **Add base entry to realm** if the LDAP repository that you need is not listed in the collection. Then complete the steps in “Adding an external repository in a federated repository configuration” on page 196.
7. On the Federated repositories panel, complete the following steps:
 - a. Optional: Repeat step 6 if the LDAP repository that you need is not listed in the collection.
 - b. Optional: Select the repositories in the collection that you do not need in the realm and click **Remove**. The following restrictions apply:
 - The realm must always contain at least one base entry; therefore, you cannot remove every entry.
 - If you plan to remove the built-in, file-based repository from the administrative realm, verify that at least one user in another member repository is a console user with administrative rights. Otherwise, you must disable security to regain access to the administrative console.
 - c. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes multiple LDAP repositories, is configured.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Configuring a single built-in, file-based repository and one or more Lightweight Directory Access Protocol repositories in a federated repository configuration

Follow this task to configure a single built-in, file-based repository and multiple Lightweight Directory Access Protocol (LDAP) repositories in a federated repository configuration.

Before you begin

To configure a built-in, file-based repository in a federated repository configuration, you must know the primary administrative user name of the user who manages WebSphere Application Server resources and user accounts.

To configure an LDAP repository in a federated repository configuration, you must know a valid user name (ID), the user password, the server host and port and, if necessary, the bind distinguished name (DN) and the bind password. You can choose any valid user in the repository that is searchable. In some LDAP servers, administrative users are not searchable and cannot be used (for example, cn=root in SecureWay). This user is referred to as a WebSphere Application Server *administrative user name* or *administrative ID* in the documentation. Being an administrative ID means a user has special privileges when calling some protected internal methods. Normally, this ID and password are used to log in to the administrative console after you turn on security. You can use other users to log in if those users are part of the administrative roles.

Note: Client certificate login is not supported in a realm that includes a single built-in, file-based repository or a single built-in, file-based repository with other repositories.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Enter the name of the realm in the Realm name field. You can change the existing realm name.
4. Enter the name of the primary administrative user in the Primary administrative user name field, for example, adminUser.

Note: When you configure multiple repositories that includes a single built-in, file-based repository, the primary administrative user name must exist in the file-based repository. If the primary administrative user name does not exist in the file-based repository, then the name is created in the file-based repository. The primary administrative user name cannot exist in other repositories.

5. Select the **Ignore case for authorization** option.

Note: When the realm includes a built-in, file-based repository, you must enable the **Ignore case for authorization** option.

When you enable this option, the authorization check is case-insensitive. Normally, an authorization check involves checking the complete DN of a user, which is unique in the realm and is case-insensitive. Clear this option when all of the member repositories in the realm are case-sensitive.

Note: Some repositories contain data that is case-sensitive only, and some repositories contain data that is case-insensitive only. Do not include both case-sensitive and case-insensitive repositories in the realm. For example, do not include case-sensitive repositories in the realm with a built-in, file-based repository.

6. Optional: Click **Add base entry to realm** if the LDAP repository that you need is not contained in the collection. Then complete the steps in “Adding an external repository in a federated repository configuration” on page 196.
7. On the Federated repositories panel, complete the following steps:
 - a. Optional: Repeat step 6 if the LDAP repository that you need is not listed in the collection.
 - b. Click **Use built-in repository** if the built-in, file-based repository is not listed in the collection.
 - c. Optional: Select the repositories in the collection that you do not need in the realm and click **Remove**.

Note: The realm must always contain at least one base entry; therefore, you cannot remove every entry.

- d. Click **OK**.
8. Provide an administrative user password. This panel displays only when a built-in, file-based repository is included in the realm. Otherwise, the panel does not display. If a built-in, file-based repository is included, complete the following steps:
 - a. Supply a password for the administrative user in the Password field.
 - b. Confirm the password of the primary administrative user in the Confirm password field.
 - c. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes a single built-in, file-based repository and one or more LDAP repositories, is configured.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Manually configuring a Lightweight Directory Access Protocol repository in a federated repository configuration

Follow this topic to manually configure Lightweight Directory Access Protocol (LDAP) repository in a federated repository configuration.

Before you begin

As a prerequisite, you need to add a LDAP repository to your WebSphere Application Server configuration, where you define the following information:

Item Name	Example
Repository identifier	ldaprepo1
Directory type	IBM Tivoli Directory Server
Primary host name	localhost
Port	389
Bind distinguished name	cn=ldapadmin
Bind password	yourpwd
Login properties	uid (a property containing login information)

See “Lightweight Directory Access Protocol repository configuration settings” on page 193 for the specific steps you must perform to establish this LDAP repository.

About this task

At this point, you have a valid LDAP repository ready to be manually configured in a federated repository configuration.

1. Map the federated repository entity types to the LDAP object classes.
 - a. Configure the LDAP repository to match the used LDAP object class for users.
 - 1) In the administrative console, click **Security > Global security**.
 - 2) Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
 - 3) Under Related items, click **Manage repositories**.
 - 4) Select the repository (for example, ldaprepo1).
 - 5) Click **LDAP entity types**.
 - 6) Click **PersonAccount**.
 - 7) Insert the *objectclass* name used in our LDAP server, for example, inetOrgPerson.

8) Click **Apply**.

9) Click **Save**.

See “Configuring supported entity types in a federated repository configuration” on page 213 for an explanation of the supported entity types.

See <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/topic/com.ibm.websphere.wim.doc.en/ldap.html> for a description of the LDAP default mappings.

b. Configure the LDAP repository to match the used LDAP *objectclass* for groups

1) In the administrative console, click **Security > Global security**.

2) Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.

3) Under Related items, click **Manage repositories**.

4) Select `ldaprepo1`.

5) Click **LDAP entity types**.

6) Click **Group**.

7) Insert the objectclass name used for your LDAP server, for example, `groupOfUniqueNames`.

8) Click **Apply**.

9) Click **Save**.

See “Group attribute definition settings” on page 231 for an explanation of group attribute definitions.

2. Map the federated repository property names to the LDAP attribute names.

a. Configure the LDAP repository to match the used LDAP attributes for a user.

1) Edit the file

```
{WAS_HOME}\profiles\{profileName}\config\cells\{cellName}\wim\config\wimconfig.xml
```

2) Look for the section in this file containing the LDAP repository configuration, For example,

```
a) <config:repositories
  xsi:type="config:LdapRepositoryType"
  adapterClassName="com.ibm.ws.wim.adapter.ldap.LdapAdapter"
  id="ldaprepo1" ...>
```

```
b) <config:attributeConfiguration>
```

```
c) ...
```

```
d) <config:attributes name="anLDAPAttribute"
  propertyName="aVMMAtribute"/>
```

```
e) ...
<config:attributeConfiguration>
```

3) Add an element of type `config:attributes` to define the mapping between a given federated depository property name, such as `departmentNumber`, to a desired LDAP attribute name, such as `warehouseSection`.

Note: For all given federated depository properties, a one-to-one mapping is assumed. If no explicit mapping of the above type is defined, for example the federated repository property `departmentNumber`, the underlying LDAP attribute name, `departmentNumber` is assumed.

b. Configure the unsupported properties of the federated repository.

To indicate that a given federated repository property, such as `departmentNumber` is not supported by any LDAP attributes, you need to define the following type of element:

```
<config:repositories xsi:type="config:LdapRepositoryType"
  adapterClassName="com.ibm.ws.wim.adapter.ldap.LdapAdapter"
  id="ldaprepo1" ...>
<config:attributeConfiguration>
```

```
...
<config:propertiesNotSupported name=" departmentNumber"/>
...
<config:attributeConfiguration>
```

- c. Configure the LDAP repository to match the used LDAP user membership attribute in the groups.
 - 1) In the administrative console, click **Security > Global security**.
 - 2) Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
 - 3) Under Related items, click **Manage repositories**.
 - 4) Select `ldaprepo1`
 - 5) Click **Group attribute definitions**.
 - 6) Click **Member attributes**.
 - 7) Check if your LDAP attributes (for example, `uniqueMember`) is specified for your LDAP objectclass (for example, `groupOfUniqueNames`).
 - If not specified, click **New** and add the pair (*objectclass / member attribute name*) that applies to your LDAP schema (for example, `uniqueMember / groupOfUniqueNames`)
 - If specified, proceed.
 - 8) Click **Apply**.
 - 9) Click **Save**.
3. Map other LDAP settings by configuring a new base entry for the new LDAP repository.
 - a. In the administrative console, click **Security > Global security**.
 - b. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
 - c. Click **Add Base Entry to Realm**.
 - d. Select `ldaprepo1`.
 - e. Specify:
 - The base entry within the federated repository realm, for example, `o=Default Organization`
 - The base entry within the LDAP repository, for example, `o=Default Organization`
 - f. Click **Apply**.
 - g. Click **Save**.

For an explanation of base entries, see “Configuring supported entity types in a federated repository configuration” on page 213

Results

After completing these steps, your federated repository matches the LDAP server settings.

What to do next

Related tasks

“Configuring supported entity types in a federated repository configuration” on page 213
Follow this task to configure supported entity types for user and group management.

Related information

“Lightweight Directory Access Protocol repository configuration settings” on page 193
Use this page to configure secure access to a Lightweight Directory Access Protocol (LDAP) repository with optional failover servers.

 <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/topic/com.ibm.websphere.wim.doc.en/ldap.html>

Configuring Lightweight Directory Access Protocol in a federated repository configuration

Follow this topic to configure Lightweight Directory Access Protocol (LDAP) settings in a federated repository configuration.

Before you begin

You have chosen among various ways to configure LDAP:

- “Configuring a single, Lightweight Directory Access Protocol repository in a new configuration under Federated repositories” on page 182
- “Changing a federated repository configuration to include a single, Lightweight Directory Access Protocol repository only” on page 183
- “Configuring multiple Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 184
- “Configuring a single built-in, file-based repository and one or more Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 185
- “Managing repositories in a federated repository configuration” on page 215

About this task

At this point, you are viewing the LDAP repository configuration page of the administrative console.

1. Enter a unique identifier for the repository in the Repository identifier field. This identifier uniquely identifies the repository within the cell, for example: LDAP1.
2. Select the type of LDAP server that is used from the Directory type list. The type of LDAP server determines the default filters that are used by WebSphere Application Server.
IBM Tivoli Directory Server users can choose either IBM Tivoli Directory Server or SecureWay as the directory type. Use the IBM Tivoli Directory Server directory type for better performance. For a list of supported LDAP servers, see “Using specific directory servers as the LDAP server” on page 133.
3. Enter the fully qualified host name of the primary LDAP server in the Primary host name field. You can enter either the IP address or the domain name system (DNS) name.
4. Enter the server port of the LDAP directory in the Port field. The host name and the port number represent the realm for this LDAP server in a mixed version nodes cell. If servers in different cells are communicating with each other using Lightweight Third Party Authentication (LTPA) tokens, these realms must match exactly in all the cells.

The default value is 389, which is not a Secure Sockets Layer (SSL) connection. Use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.

If multiple WebSphere Application Servers are installed and configured to run in the same single sign-on domain, or if WebSphere Application Server interoperates with a previous version of WebSphere Application Server, then it is important that the port number match all configurations. For

example, if the LDAP port is explicitly specified as 389 in a Version 5.x or 6.0.x configuration, and WebSphere Application Server at Version 6.1 is going to interoperate with the Version 5.x or 6.0.x server, then verify that port 389 is specified explicitly for the Version 6.1 server.

5. Optional: Enter the host name of the failover LDAP server in the Failover host name field. You can specify a secondary directory server to be used in the event that your primary directory server becomes unavailable. After switching to a secondary directory server, LDAP repository attempts to reconnect to the primary directory server every 15 minutes.
6. Optional: Enter the port of the failover LDAP server in the Port field and click **Add**. The default value is 389, which is not a Secure Sockets Layer (SSL) connection. Use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.
7. Optional: Select the type of *referral*. A referral is an entity that is used to redirect a client request to another LDAP server. A referral contains the names and locations of other objects. It is sent by the server to indicate that the information that the client requested can be found at another location, possibly at another server or several servers. The default value is ignore.

ignore

Referrals are ignored.

follow Referrals are followed automatically.

8. Optional: Enter the bind DN name in the Bind distinguished name field, for example, cn=root. The bind DN is required if anonymous binds are not possible on the LDAP server to obtain user and group information or for write operations. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed. If the LDAP server is set up to use anonymous binds, leave this field blank. If a name is not specified, the application server binds anonymously.

Note: To create LDAP queries or to browse, an LDAP client must bind to the LDAP server using the distinguished name (DN) of an account that has the authority to search and read the values of LDAP attributes, such as user and group information. The LDAP administrator ensures that **read access privileges** are set for the bind DN. Read access privileges allow access to the subtree of the base DN and ensure that searches of user and group information are successful.

The directory server provides an operational attribute in each directory entry (for example, the IBM Directory Server uses `ibm-entryUuid` as the operational attribute). The value of this attribute is a universally unique identifier (UUID), which is chosen automatically by the directory server when the entry is added, and is expected to be unique: no other entry with the same or different name would have this same value. Directory clients may use this attribute to distinguish objects identified by a distinguished name or to locate an object after renaming.

Ensure that the bind credentials have the authority to read this attribute.

9. Optional: Enter the password that corresponds to the bind DN in the Bind password field.
10. Optional: Enter the property names to use to log into WebSphere Application Server in the Login properties field. This field takes multiple login properties, delimited by a semicolon (;). For example, `uid;mail`.
All login properties are searched during login. If multiple entries or no entries are found, an exception is thrown. For example, if you specify the login properties as `uid;mail` and the login ID as Bob, the search filter searches for `uid=Bob` or `mail=Bob`. When the search returns a single entry, then authentication can proceed. Otherwise, an exception is thrown.
11. Optional: Select the certificate map mode in the Certificate mapping field. You can use the X.590 certificates for user authentication when LDAP is selected as the repository. The Certificate mapping field is used to indicate whether to map the X.509 certificates into an LDAP directory user by `EXACT_DN` or `CERTIFICATE_FILTER`. If `EXACT_DN` is selected, the DN in the certificate must exactly match the user entry in the LDAP server, including case and spaces.

12. If you select **CERTIFICATE_FILTER** in the Certificate mapping field, specify the LDAP filter for mapping attributes in the client certificate to entries in LDAP.

If more than one LDAP entry matches the filter specification at run time, authentication fails because the result is an ambiguous match. The syntax or structure of this filter is:

```
LDAP attribute=${Client certificate attribute}
```

For example, uid=\${SubjectCN}.

The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. The right side must begin with a dollar sign (\$) and open bracket ({} and end with a close bracket (}). You can use the following certificate attribute values on the right side of the filter specification. The case of the strings is important:

- \${UniqueKey}
- \${PublicKey}
- \${PublicKey}
- \${Issuer}
- \${NotAfter}
- \${NotBefore}
- \${SerialNumber}
- \${SigAlgName}
- \${SigAlgOID}
- \${SigAlgParams}
- \${SubjectCN}
- \${Version}

13. Optional: Select the **Require SSL communications** option if you want to use Secure Sockets Layer communications with the LDAP server.

If you select the **Require SSL communications** option, you can select either the **Centrally managed** or **Use specific SSL alias** option.

Centrally managed

Enables you to specify an SSL configuration for a particular scope, such as the cell, node, server, or cluster in one location. To use the Centrally managed option, you must specify the SSL configuration for the particular set of endpoints. The Manage endpoint security configurations and trust zones panel displays all of the inbound and outbound endpoints that use the SSL protocol. If you expand the Inbound or Outbound section of the panel and click the name of a node, you can specify an SSL configuration that is used for every endpoint on that node. For an LDAP registry, you can override the inherited SSL configuration by specifying an SSL configuration for LDAP. To specify an SSL configuration for LDAP, complete the following steps:

- a. Click **Security > SSL certificate and key management > Manage endpoint security configurations and trust zones**.
- b. Expand **Outbound > cell_name > Nodes > node_name > Servers > server_name > LDAP**.

Use specific SSL alias

Select the **Use specific SSL alias** option if you intend to select one of the SSL configurations in the menu that follows the option.

This configuration is used only when SSL is enabled for LDAP. The default is *DefaultSSLSettings*. To modify or create a new SSL configuration, complete the following steps:

- a. Click **Security > SSL certificate and key management**.
- b. Under Configuration settings, click **Manage endpoint security configurations and trust zones > configuration_name**.
- c. Under Related items, click **SSL configurations**.

14. Click **OK**.

Results

After completing these steps, your LDAP repository settings are configured.

What to do next

Return to the appropriate task to complete the steps for your federated repository configuration:

- “Configuring a single, Lightweight Directory Access Protocol repository in a new configuration under Federated repositories” on page 182
- “Changing a federated repository configuration to include a single, Lightweight Directory Access Protocol repository only” on page 183
- “Configuring multiple Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 184
- “Configuring a single built-in, file-based repository and one or more Lightweight Directory Access Protocol repositories in a federated repository configuration” on page 185
- “Managing repositories in a federated repository configuration” on page 215

Lightweight Directory Access Protocol repository configuration settings:

Use this page to configure secure access to a Lightweight Directory Access Protocol (LDAP) repository with optional failover servers.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Configuring a property extension repository in a federated repository configuration” on page 197

Follow this task to configure a property extension repository to store attributes that cannot be stored in your Lightweight Directory Access Protocol (LDAP) server.

“Managing the realm in a federated repository configuration” on page 170

Follow this topic to manage the realm in a federated repository configuration.

Repository identifier:

Specifies a unique identifier for the LDAP repository. This identifier uniquely identifies the repository within the cell, for example: LDAP1.

Directory type:

Specifies the type of LDAP server to which you connect.

Expand the drop-down list to display a list of LDAP directory types.

Primary host name:

Specifies the host name of the primary LDAP server. This host name is either an IP address or a domain name service (DNS) name.

Port:

Specifies the LDAP server port.

The default value is 389, which is not a Secure Sockets Layer (SSL) connection. Use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.

Data type:	Integer	
Default:	389	
Range:	389, which is not a Secure Sockets Layer (SSL) connection	636, which is a Secure Sockets Layer (SSL) connection

Failover host name:

Specifies the host name of the failover LDAP server.

You can specify a secondary directory server to be used in the event that your primary directory server becomes unavailable. After switching to a secondary directory server, the LDAP repository attempts to reconnect to the primary directory server every 15 minutes.

Port:

Specifies the port of the failover LDAP server.

The default value is 389, which is not a Secure Sockets Layer (SSL) connection. Use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.

Data type:	Integer	
Range:	389, which is not a Secure Sockets Layer (SSL) connection	636, which is a Secure Sockets Layer (SSL) connection

Support referrals to other LDAP servers:

Specifies how referrals that are encountered by the LDAP server are handled.

A referral is an entity that is used to redirect a client request to another LDAP server. A referral contains the names and locations of other objects. It is sent by the server to indicate that the information that the client requested can be found at another location, possibly at another server or several servers. The default value is ignore.

Default:	ignore
Range:	ignore Referrals are ignored. follow Referrals are followed automatically.

Bind distinguished name:

Specifies the distinguished name (DN) for the application server to use when binding to the LDAP repository.

If no name is specified, the application server binds anonymously. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed.

Bind password:

Specifies the password for the application server to use when binding to the LDAP repository.

Login properties:

Specifies the property names to use to log into the application server.

This field takes multiple login properties, delimited by a semicolon (;). For example, `uid;mail`. All login properties are searched during login. If multiple entries or no entries are found, an exception is thrown. For example, if you specify the login properties as `uid;mail` and the login ID as Bob, the search filter searches for `uid=Bob` or `mail=Bob`. When the search returns a single entry, then authentication can proceed. Otherwise, an exception is thrown.

Certificate mapping:

Specifies whether to map X.509 certificates into an LDAP directory by `EXACT_DN` or `CERTIFICATE_FILTER`. Specify `CERTIFICATE_FILTER` to use the specified certificate filter for the mapping.

Certificate filter:

Specifies the filter certificate mapping property for the LDAP filter. The filter is used to map attributes in the client certificate to entries in the LDAP repository.

If more than one LDAP entry matches the filter specification at run time, authentication fails because the result is an ambiguous match. The syntax or structure of this filter is:

```
LDAP attribute=${Client certificate attribute}
```

For example, `uid=${SubjectCN}`.

The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. The right side must begin with a dollar sign (\$) and open bracket ({} and end with a close bracket (}). You can use the following certificate attribute values on the right side of the filter specification. The case of the strings is important:

- `${UniqueKey}`
- `${PublicKey}`
- `$(PublicKey}`
- `$(Issuer}`
- `$(NotAfter}`
- `$(NotBefore}`
- `$(SerialNumber}`
- `$(SigAlgName}`
- `$(SigAlgOID}`
- `$(SigAlgParams}`
- `$(SubjectCN}`
- `$(Version}`

Require SSL communications:

Specifies whether secure socket communication is enabled to the LDAP server.

When enabled, the Secure Sockets Layer (SSL) settings for LDAP are used, if specified.

Centrally managed:

Specifies that the selection of an SSL configuration is based upon the outbound topology view for the Java Naming and Directory Interface (JNDI) platform.

Centrally managed configurations support one location to maintain SSL configurations, rather than spreading them across the configuration documents.

Default: Enabled
Range: Enabled or Disabled

Use specific SSL alias:

Specifies the SSL configuration alias to use for LDAP outbound SSL communications.

This option overrides the centrally managed configuration for the JNDI platform.

Adding an external repository in a federated repository configuration

Follow this task to add an external repository into a federated repository configuration.

1. If the Lightweight Directory Access Protocol (LDAP) repository that you want to add to your federated repository configuration is previously configured, select the corresponding Repository on the Repository reference panel. To access the Repository reference panel, complete the following steps:
 - a. Click **Security > Global security**.
 - b. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
 - c. Click **Add base entry to realm**.

2. Enter a distinguished name for the realm base entry in the Distinguished name that uniquely identifies... field. This base entry must uniquely identify the external repository in the realm. If multiple repositories are included in the realm, use this field to define an additional distinguished name (DN) that uniquely identifies this set of entries within the realm. For example, repositories LDAP1 and LDAP2 might both use o=ibm,c=us as the base entry in the repository. Use the DN in this field to uniquely identify this set of entries in the realm. For example: o=ibm,c=us for LDAP1 and o=ibm2,c=us for LDAP2. The specified DN in this field maps to the LDAP DN of the base entry within the repository.

3. Enter the LDAP DN of the base entry within the repository in the Distinguished name of a base entry... field. The base entry indicates the starting point for searches in this LDAP directory server. This entry and its descendents are mapped to the subtree that is identified by this unique base name entry field. For example, for a user with a DN of cn=John Doe, ou=Rochester, o=IBM, c=US, specify the LDAP base entry as any of the following options:

ou=Rochester, o=IBM, c=us or o=IBM, c=us or c=us

In most cases, this LDAP DN is the same as the distinguished name for the realm base entry.

If this field is left blank, then the subtree defaults to the root of the LDAP repository. Consult your LDAP administrator to determine if your LDAP repository provides support to search from the root, or create users and groups under the root without defining a suffix beforehand.

In WebSphere Application Server, the distinguished name is normalized according to the LDAP specification. Normalization consists of removing spaces in the base distinguished name before or after commas and equal symbols. An example of a non-normalized base distinguished name is o = ibm, c = us or o=ibm, c=us. An example of a normalized base distinguished name is o=ibm,c=us.

4. If the LDAP repository that you want to add to your realm is not previously configured, complete the following steps:
 - a. Click **Add Repository** on the Repository reference panel to configure the LDAP repository. See step 1 to access the Repository reference panel.

- b. Configure LDAP on the LDAP configuration panel, as described in “Configuring Lightweight Directory Access Protocol in a federated repository configuration” on page 190.
 - c. Select the new Repository on the Repository reference panel.
5. Click **OK**.

Results

You have added a new or previously configured external repository into your federated repository configuration.

What to do next

1. Before you can manage this account with Users and Groups, configure supported entity types as described in “Configuring supported entity types in a federated repository configuration” on page 213.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Configuring a property extension repository in a federated repository configuration

Follow this task to configure a property extension repository to store attributes that cannot be stored in your Lightweight Directory Access Protocol (LDAP) server.

About this task

For security and business reasons, you might not want to not allow write operations to your repositories. However, applications calling the federated repository configuration might need to store additional properties for the entities. A federated repository configuration provides a *property extension repository*, which is a database regardless of the type of main profile repositories, for a propertylevel join configuration. For example, a company that uses an LDAP directory for its internal employees and a database for external customers and business partners might not allow write access to its LDAP and its database. The company can use the *property extension repository* in a federated repository configuration to store additional properties for the people in those repositories, excluding the user ID. When an application uses the federated repository configuration to retrieve an entry for a person, the federated repository configuration transparently joins the properties of the person that is retrieved from either the LDAP or the customer’s database with the properties of the person that is retrieved from the property extension repository into a single logical person entry.

When you configure a property extension repository, you can supply a valid data source, a direct connection configuration, or both. The system first tries to connect by way of the data source. If the data source is not available, then the system uses the direct access configuration.

Note: You cannot configure a property extension repository in a mixedversion deployment manager cell.

1. Configure the WebSphere Application Server data source. See “Configuring the WebSphere Application Server data source” on page 209.
2. If you are adding new properties (including properties that are stored in the property extension repository) to the schema, you must do the following before you create the property extension repository.

- a. Open or create the `wimxmlextension.xml` file under the `profile_root\config\cells\cell_name\wim\model` directory.

Note: Make sure the editor is on the deployment manager node.

- b. Add the schema definition of the new property. The following sample `wimxmlextension.xml` file adds a new property called `ibmotherEmail` to both the `Person` and `PersonAccount` entity types. This new property type is "String" and it is multiplevalued.

```
<sdo:datagraph xmlns:sdo="commonj.sdo"
  xmlns:wim="http://www.ibm.com/websphere/wim">
  <wim:schema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="true"
      propertyName="ibm-otherEmail">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="true"
      propertyName="ibm-personalTitle">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="true"
      propertyName="ibm-middleName">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="true"
      propertyName="ibm-generationQualifier">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-regionalLocale">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-timeZone">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-preferredCalendar">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-alternativeCalendar">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-firstDayOfWeek">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-firstWorkDayOfWeek">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="false"
      propertyName="ibm-gender">
    <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
    <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
```

```

    <wim:propertySchema nsURI="http://www.ibm.com/websphere/wim" dataType="String" multiValued="true"
      propertyName="ibm-hobby">
      <wim:applicableEntityTypeNames>Person</wim:applicableEntityTypeNames>
      <wim:applicableEntityTypeNames>PersonAccount</wim:applicableEntityTypeNames>
    </wim:propertySchema>
  </wim:schema>
</sdo:datagraph>

```

Available data types are defined in `com.ibm.websphere.wim.SchemaConstants`. For example:

```

/**
 * Instance Class: java.lang.String
 */
String DATA_TYPE_STRING = "String";
/**
 * Instance Class: int
 */
String DATA_TYPE_INT = "Int";
/**
 * Instance Class: java.lang.Object
 */
String DATA_TYPE_DATE = "Date";
/**
 * Instance Class: dobjava.lang.Object
 */
String DATA_TYPE_ANY_SIMPLE_TYPE = "AnySimpleType";
/**
 * Instance Class: java.lang.String
 */
String DATA_TYPE_ANY_URI = "AnyURI";
/**
 * Instance Class: java.lang.boolean
 */
String DATA_TYPE_BOOLEAN = "Boolean";
/**
 * Instance Class: long
 */
String DATA_TYPE_LONG = "Long";
/**
 * Instance Class: double
 */
String DATA_TYPE_DOUBLE = "Double";
/**
 * Instance Class: short
 */
String DATA_TYPE_SHORT = "Short";

```

- c. Follow the example inside this file to define the new property definitions. The schema file for `wimlaproperties.xml` is `wimdbproperty.xsd` and is in the same directory. It can be used for reference.
3. Run the `setupIdMgrPropertyExtensionRepositoryTables` command to create the property extension repository and to add the new properties.
4. Set up the property extension repository using `wsadmin` by following the procedure discussed in “Setting up an entry mapping repository, a property extension repository, or a custom registry database repository using `wsadmin` commands” on page 202; ignore the “Before you begin” options.
5. Configure the property extension repository by completing the following steps:
 - a. In the administrative console, click **Security > Global security**.
 - b. Under User account repository, select **Federated repositories**, and click **Configure**.
 - c. Click **Property extension repository**.
 - d. Supply the name of the data source in the Data source name field.
 - e. Select the type of database that is used for the property extension repository.
 - f. Supply the name of the Java database connectivity (JDBC) driver in the JDBC driver field.

Values include:

DB2® com.ibm.db2.jcc.DB2Driver

Oracle

oracle.jdbc.driver.OracleDriver

Informix®

com.informix.jdbc.IfxDriver

Microsoft SQL Server

com.microsoft.jdbc.sqlserver.SQLServerDriver

Derby org.apache.derby.jdbc.EmbeddedDriver

- g. Supply the database URL that is used to access the property extension repository with JDBC in the Database URL field. Use an alphanumeric text string that conforms to the standard JDBC URL syntax.

Values include:

DB2 jdbc:db2://<hostname>:<port>/<DB2location>

Oracle

jdbc:oracle:thin:@<hostname>:<port>:<dbname>

Derby jdbc:derby:c:\derby\wim

Microsoft SQL Server

jdbc:microsoft:sqlserver://<hostname>:1433;databaseName=wim;selectmethod=cursor;

Informix

jdbc:informixqli://<hostname>:1526/wim:INFORMIXSERVER=<IFXServerName>;

- h. Supply the user name of the database administrator in the Database administrator user name field.
- i. Supply the password of the database administrator in the Password field.
- j. Specify the entity retrieval limit in the Entity retrieval limit field. The entity retrieval limit is the maximum number of entities that the system can retrieve from the property extension repository with a single database query. The default value is 200.
- k. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes a property extension repository, is configured.

What to do next

1. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Secure administration, applications, and infrastructure panel.
2. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Property extension repository settings:

Use this page to configure a property extension repository that is used to store attributes that cannot be stored in existing repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.

3. Under Additional properties, click **Property extension repository**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Data source name:

Specifies the Java Naming and Directory Interface (JNDI) name of the data source that is used to access the property extension repository.

Default: jdbc/wimDS

Database type:

Specifies the type of database that is used for the property extension repository.

Default: DB2

JDBC driver:

Specifies the Java Database Connectivity (JDBC) driver that is used to access the entry mapping repository.

Values include:

DB2 COM.ibm.db2.jcc.DB2Driver

Oracle
oracle.jdbc.driver.OracleDriver

Informix
com.informix.jdbc.IfxDriver

DataDirect Connect
com.ddtek.jdbc.sqlserver.SQLServerDriver

Derby org.apache.derby.jdbc.EmbeddedDriver

Microsoft SQL Server
com.microsoft.sqlserver.jdbc.SQLServerDriver

Database URL:

Specifies the Web address for the property extension repository.

Values include:

DB2 jdbc:db2:wim

Informix
jdbc:informix-sqli://host_name:1526/wim:INFORMIXSERVER=IFXServerName;

DataDirect Connect
jdbc:datadirect:sqlserver://host_name:1433;databaseName=wim;selectMethod=cursor;

Derby jdbc:derby:c:\derby\wim

Oracle

```
jdbc:oracle:thin:@host_name:port:dbname
```

Microsoft SQL Server

```
jdbc:sqlserver://host_name:1433;databaseName=wim;selectMethod=cursor;
```

Database administrator user name:

Specifies the user name of the database administrator that is used to access the property extension repository.

Password:

Specifies the password that is used to enable the database administrator to access the property extension repository.

Entity retrieval limit:

Specifies the maximum number of entities that the system can retrieve from the property extension repository with a single database query.

Data type:	Integer
Default:	200

Setting up an entry mapping repository, a property extension repository, or a custom registry database repository using wsadmin commands:

You can set up an entry mapping repository, a property extension repository, or a custom registry database repository using wsadmin commands.

Before you begin

If you are setting up an entry mapping repository, begin with the steps described in “Configuring an entry mapping repository in a federated repository configuration” on page 210.

If you are setting up a property extension repository, begin with the steps described in “Configuring a property extension repository in a federated repository configuration” on page 197.

About this task

When you create a repository, use the appropriate wsadmin commands to define the database schema and to populate the database property definitions.

1. Create the database. You can use any relational database product. The following examples give you tips for specific vendors.
 - a. For **DB2**, open a DB2 command window or command center and enter the following:

```
db2 create database <name> using codeset UTF-8 territory US
```

Enter the following database tuning commands:

```
db2 update database configuration for <name> using applheapsz 1024
db2 update database configuration for <name> using stmtheap 4096
db2 update database configuration for <name> using app_ctl_heap_sz 2048
db2 update database configuration for <name> using locklist 1024
db2 update database configuration for <name> using indexrec RESTART
db2 update database configuration for <name> using logfilsiz 1000
db2 update database configuration for <name> using logprimary 12
db2 update database configuration for <name> using logsecond 10
db2 update database configuration for <name> using sortheap 2048
db2set DB2_RR_TO_RS=yes
```

- b. Optional: For **Informix** databases using dbaccess, enter the following command:

```
CREATE DATABASE <name> WITH BUFFERED LOG
```
 - c. Optional: For **Oracle** databases, the database should already exist during Oracle installation (for example, orcl).
2. Run the `setupIdMgrEntryMappingRepositoryTables` command, the `setupIdMgrPropertyExtensionRepositoryTables` command, or the `setupIdMgrDBTables` command (for custom registry repositories) by doing the following:
 - a. Start WebSphere Application Server.
 - b. Open a command window and go to the `<WAS>/Profiles/<PROFILE_NAME>bin` directory.
 - c. Start `wsadmin`.
 - d. Type the necessary commands as described below.

What to do next

Using these commands, you can:

- Specify the arguments on the command line.
- Specify the arguments in a file.

The `-file` option enables you to specify a file in which some or all of the parameters are specified. To use the `-file` argument on the command line, enter the full path to the file. Parameters in the file must be specified in `key=value` pairs and each must be on its own line. If a parameter is specified on both the command line and in the file, the value on the command line takes precedence.

Tips for diagnosing argument errors:

- If an argument is not properly specified on the command line or in the file, a message is returned which states that the argument was not properly specified. This might mean that the argument was not specified at all or was required for a given configuration but was not specified.
- If the argument was not specified at all, check that the parameter is specified on the command line or in the file, and that it is properly spelled and has matching case.
- If the argument was required for a given configuration but was not specified, it is possible that a value is not required solely by the command but is required for the type of database and configuration you are setting.

For example, if you set the `dn`, `wasAdminId`, or `wasAdminPassword` parameters, you must also specify the `dbDriver` parameter.

Additionally, if the `dn`, `wasAdminId` or `wasAdminPassword` parameters are specified, and the `databaseType` is not a Apache Derby v10.2 database, then the `dbAdminId` and `dbAdminPassword` parameters must also be specified.

The `setupIdMgrDBTables` command:

The `setupIdMgrDBTables` command creates, and populates the tables in the database that you previously created. Required arguments are prefixed by a double start (**). Arguments are case-sensitive, both through the command line and the file.

Parameters:

**`schemaLocation` (String)

The location of the `<WAS>/etc/wim/setup` directory.

`dbPropXML` (String)

The location of database repository property definition XML file.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbDriver (String)

The name of the database driver. For example: com.ibm.db2.jcc.DB2Driver.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

dn (String)

The default organization uniqueName to replace. For example: o=yourco. If it is not set, o=DefaultOrganization is used.

wasAdminId (String)

The WebSphere Application Server admin user ID. The ID should be a short name, not a uniqueName. For example: wasadmin. After creation, the uniqueName is uid=wasadmin, <defaultOrg>.

wasAdminPassword (String)

The WebSphere Application Server admin user password. If wasAdminId is set, then this parameter is mandatory.

saltLength (Integer)

The salt length of the randomly generated salt for password hashing.

encryptionKey (String)

The password encryption key. Set the password encryption key to match the encryption key in the wimconfig.xml file for the repository. If the encryption key is not set, the default is used.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

The deleteIdMgrDBTables command:

The deleteIdMgrDBTables command deletes the tables in the database.

Parameters:

****schemaLocation (String)**

The location of the <WAS>/etc/wim/setup directory.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbDriver (String)

The name of the database driver. For example: com.ibm.db2.jcc.DB2Driver.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

The setupIdMgrPropertyExtensionRepositoryTables command:

The setupIdMgrPropertyExtensionRepositoryTables command sets up the property extension repository, which includes creating and populating the tables in the database.

*Parameters:*****schemaLocation (String)**

The location of the <WAS>/etc/wim/setup directory.

laPropXML (String)

The location of the property extension repository definition XML file.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

The deleteIdMgrPropertyExtensionRepositoryTables command:

The deleteIdMgrPropertyExtensionRepositoryTables command deletes the tables in the property extension database.

Parameters:

****schemaLocation (String)**

The location of the <WAS>/etc/wim/setup directory.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbDriver (String)

The name of the database driver. For example: com.ibm.db2.jcc.DB2Driver.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

The setupIdMgrEntryMappingRepositoryTables command:

The setupIdMgrEntryMappingRepositoryTables command sets up the entry mapping repository, which includes creating and populating the tables of the repository.

Parameters:

****schemaLocation (String)**

The location of the <WAS>/etc/wim/setup directory.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbDriver (String)

The name of the database driver. For example: com.ibm.db2.jcc.DB2Driver.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

The deleteIdMgrEntryMappingRepositoryTables command:

The deleteIdMgrEntryMappingRepositoryTables command deletes the tables in the entry mapping repository.

*Parameters:*****schemaLocation (String)**

The location of the <WAS>/etc/wim/setup directory.

****databaseType (String)**

The type of database. Supported databases are db2, oracle, informix, derby, sqlserver, db2zos, and db2iseries.

****dbURL (String)**

The database URL for direct access mode. For example: jdbc:db2:wim.

dbDriver (String)

The name of the database driver. For example: com.ibm.db2.jcc.DB2Driver.

dbAdminId (String)

The database administrator ID for direct access mode. For example: db2admin.

Note: For a Apache Derby v10.2 embedded database, dbAdminId is not required.

dbAdminPassword (String)

The password associated with the dbAdminId.

Note: For a Apache Derby v10.2 embedded database, dbAdminPassword is not required.

derbySystemHome (String)

The home location of the Apache Derby v10.2 system if you are setting up a Apache Derby v10.2 database.

reportSqlError (String)

Specifies whether to report SQL errors while setting up databases.

file (String)

The full path to a file containing the input parameters. Each input parameter must match a corresponding parameter as it would be typed on the command line, and it must be placed in a key=value pair. Each pair must be on a separate line.

Sample command line usage:

To set up a database using the command line, enter the following:

```
$AdminTask setupIdMgrDBTables {-schemaLocation "C:\WAS7\etc\wim\setup" -dbPropXML  
"C:\WAS7\etc\wim\setup\wimdbproperties.xml" -databaseType db2  
-dbURL jdbc:db2:wim -dbAdminId db2admin  
-dbDriver com.ibm.db2.jcc.DB2Driver -dbAdminPassword db2adminPwd  
-reportSqlError true}
```

To delete database tables using the command line, enter the following:

```
$AdminTask deleteIdMgrDBTables {-schemaLocation "C:\WAS7\etc\wim\setup"  
-databaseType db2 -dbURL jdbc:db2:wim -dbAdminId db2admin  
-dbDriver com.ibm.db2.jcc.DB2Driver -dbAdminPassword db2adminPwd  
-reportSqlError true}
```

To set up a property extension repository using the command line, enter the following:

```
$AdminTask setupIdMgrPropertyExtensionRepositoryTables {-schemaLocation  
"C:\WAS7\etc\wim\setup"  
-taPropXML "C:\WAS7\etc\wim\setup\wim1aproperties.xml" -databaseType db2  
-dbURL jdbc:db2:wim -dbAdminId db2admin -dbDriver com.ibm.db2.jcc.DB2Driver  
-dbAdminPassword db2adminPwd -reportSqlError true}
```

To delete a property extension repository using the command line, enter the following:

```
$AdminTask deleteIdMgrPropertyExtensionRepositoryTables {-schemaLocation "C:\WAS7\etc\wim\setup "  
-databaseType db2 -dbURL jdbc:db2:wim -dbAdminId db2admin -dbDriver  
com.ibm.db2.jcc.DB2Driver -dbAdminPassword db2adminPwd -reportSqlError true}
```

To set up an entry mapping repository using the command line, enter the following:

```
$AdminTask setupIdMgrEntryMappingRepositoryTables {-schemaLocation "C:\WAS7\etc\wim\setup"  
-databaseType db2 -dbURL jdbc:db2:wim -dbAdminId db2admin -dbDriver  
com.ibm.db2.jcc.DB2Driver -dbAdminPassword db2adminPwd -reportSqlError true}
```

To delete an entry mapping repository using the command line, enter the following:

```
$AdminTask deleteIdMgrEntryMappingRepositoryTables {-schemaLocation "C:\WAS7\etc\wim\setup"  
-databaseType db2 -dbURL jdbc:db2:wim -dbAdminId db2admin -dbDriver  
com.ibm.db2.jcc.DB2Driver -dbAdminPassword db2adminPwd -reportSqlError true}
```

Sample CLI Usage using -file option:

To set up a database with the -file option using the example params.txt file below, enter the following:

```
$AdminTask setupIdMgrDBTables {-file C:\params.txt -dbPropXML  
"C:\OverrideDBPropParam\wimdbproperties.xml"}
```

Params.txt

```

schemaLocation=C:\WAS7\etc\wim\setup
dbPropXML=C:\Program Files\IBM\WebSphere\AppServer\profiles\default
\config\cells\mycell\wim\config\wimdbproperties.xml
IaPropXML=C:\Program Files\IBM\WebSphere\AppServer\profiles\default
\config\cells\mycell\wim\config\wimIaproperties.xml
databaseType=db2
dbURL=jdbc:db2:wim
dbDriver=com.ibm.db2.jcc.DB2Driver
reportSqlError=true
dn=o=db.com
dbAdminId=db2admin
dbAdminPassword=dbPassword
wasAdminId=wasadmin
wasAdminPassword=wasadmin1

```

To set up a database with the `-file` option using a file only, enter the following:

```
$AdminTask setupIdMgrDBTables {-file C:\params.txt}
```

Note: The use of a file only works if `-file` is the only parameter specified on the command line. If other parameters are specified then the file is completely ignored, and only the parameters on the command line are used to execute the command.

Configuring the WebSphere Application Server data source:

Installed applications use data sources as resources to obtain connection to relational databases. To create these connections between an application and a relational database, WebSphere Application Server uses the driver implementation classes that are encapsulated by the JDBC provider, which is an object that represents vendor-specific JDBC driver classes to WebSphere Application Server. For access to a relational databases, applications use the JDBC drivers and data sources that you configure for WebSphere Application Server.

1. Start the WebSphere Application Server administrative console.
2. Click **Security -> Global security**.
3. On the Configuration panel, expand Java Authentication and Authorization Service and click **J2C authentication data**.
4. Click **New** and enter the Alias, User ID and Password.
5. Click **Ok**.
6. On the WebSphere Application Server administrative console, expand **Resources**. Expand **JDBC** then click **JDBC Providers**.
7. In the Scope section, choose the **Node level** from the drop-down list.
8. Click **New** to create a new JDBC driver.
9. Select, in this order, the Database type, Provider type, Implementation type and Name. The Name automatically fills based on the implementation type you choose.
10. Click **Next** and configure the database class path. Click **Next**.
11. On the Summary page, click **Finish**.
12. Click **Save** to save your selections. The JDBC providers page then appears.
13. On the WebSphere Application Server administrative console, click **Data sources**.
14. Click **New** to create a new data source. Enter the Data source name and the JNDI name, and choose the authentication alias from the drop-down list in Component-managed authentication alias. The JNDI name should match the `datasourceName` value set in `wimconfig.xml`. By default, it is `jdbc/wimDS`.

Note: For Apache Derby v10.2 embedded databases, leave the Component-managed authentication alias field set to NONE.

15. Click **Next**.

16. Enter the Database name and deselect the checkbox, **Use this data source in container managed persistence (CMP)**. Click **Next**.
17. On the Summary page, click **Finish**.
18. The Data sources page displays. Click **Save**. Then select the check box for the authentication alias previously created. Click **Test Connection**. The message should indicate that the connection is successful. Ignore any warnings, and then click **Next**.
19. Save the configurations, and restart WebSphere Application Server.

Configuring an entry mapping repository in a federated repository configuration

Follow this task to configure an entry mapping repository that is used to store data for managing profiles on multiple repositories.

About this task

An *entry-level join* means that the federated repository configuration uses multiple repositories simultaneously and recognizes the entries in the different repositories as entries representing distinct entities. For example, a company might have a Lightweight Directory Access Protocol (LDAP) directory that contains entries for its employees and a database that contains entries for business partners and customers. By configuring an entry mapping repository, a federated repository configuration can use both the LDAP and the database at the same time. The federated repository configuration hierarchy and constraints for identifiers provide the aggregated namespace for both of those repositories and prevent identifiers from colliding.

When you configure an entry mapping repository, you can supply a valid data source, a direct connection configuration, or both. The system first tries to connect by way of the data source. If the data source is not available, then the system uses the direct access configuration.

Note: You cannot configure an entry mapping repository in a mixed-version deployment manager cell.

1. Configure the WebSphere Application Server data source. See “Configuring the WebSphere Application Server data source” on page 209.
2. Set up the entry mapping repository using wsadmin. See “Setting up an entry mapping repository, a property extension repository, or a custom registry database repository using wsadmin commands” on page 202; ignore the “Before you begin” options.
3. Configure the entry mapping repository into the federated repository by doing the following:
 - a. In the administrative console, click **Security > Global security**.
 - b. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
 - c. Click **Entry mapping repository**.
 - d. Supply the name of the data source in the Data source name field.
 - e. Select the type of database that is used for the property extension repository.
 - f. Supply the name of the Java database connectivity (JDBC) driver in the JDBC driver field.

Values include:

DB2 `com.ibm.db2.jcc.DB2Driver`

Informix
 `com.informix.jdbc.IfxDriver`

DataDirect Connect
 `com.ddtek.jdbc.sqlserver.SQLServerDriver`

Derby `org.apache.derby.jdbc.EmbeddedDriver`

Microsoft SQL Server
 `com.microsoft.sqlserver.jdbc.SQLServerDriver`

Oracle

`oracle.jdbc.driver.OracleDriver`

- g. Supply the database URL that is used to access the property extension repository with JDBC in the Database URL field. Use an alphanumeric text string that conforms to the standard JDBC URL syntax.

Values include:

DB2 `jdbc:db2:wim`

Informix

`jdbc:informix-sqli://host_name:1526/wim:INFORMIXSERVER=IFXServerName;`

DataDirect Connect

`jdbc:datadirect:sqlserver://host_name:1433;databaseName=wim;selectMethod=cursor;`

Derby `jdbc:derby:c:\derby\wim`

Microsoft SQL Server

`jdbc:sqlserver://host_name:1433;databaseName=wim;selectMethod=cursor;`

Oracle

`jdbc:oracle:thin:@host_name:port:dbname`

- h. Supply the user name of the database administrator in the Database administrator user name field.
- i. Supply the password of the database administrator in the Password field.
- j. Click **OK**.

Results

After completing these steps, your federated repository configuration, which includes an entry mapping repository, is configured.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Entry mapping repository settings:

Use this page to configure an entry mapping repository that is used to store data for managing profiles on multiple repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Additional properties, click **Entry mapping repository**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Data source name:

Specifies the Java Naming and Directory Interface (JNDI) name of the data source that is used to access the entry mapping repository.

Default: jdbc/wimDS

Database type:

Specifies the type of database that is used to access the entry mapping repository.

Default: DB2

JDBC driver:

Specifies the Java Database Connectivity (JDBC) driver that is used to access the entry mapping repository.

Values include:

DB2 com.ibm.db2.jcc.DB2Driver

DataDirect Connect
com.ddtek.jdbc.sqlserver.SQLServerDriver

Informix
com.informix.jdbc.IfxDriver

Oracle
oracle.jdbc.driver.OracleDriver

Microsoft SQL Server
com.microsoft.sqlserver.jdbc.SQLServerDriver

Derby org.apache.derby.jdbc.EmbeddedDriver

Database URL:

Specifies the Web address for the entry mapping repository.

Values include:

DB2 jdbc:db2:wim

Derby jdbc:derby:c:\derby\wim

DataDirect Connect
datadirect:sqlserver://:host_name1433;databaseName=wim;selectMethod=cursor;

Oracle
jdbc:oracle:thin:@host_name:port:dbname

Microsoft SQL Server
jdbc:sqlserver://host_name:1433;databaseName=wim;selectMethod=cursor;

Informix
jdbc:informix-sqli://host_name:1526/wim:INFORMIXSERVER=IFXServerName;

Database administrator user name:

Specifies the user name of the database administrator that is used to access the entry mapping repository.

Password:

Specifies the password that is used to enable the database administrator to access the entry mapping repository.

Configuring supported entity types in a federated repository configuration

Follow this task to configure supported entity types for user and group management.

About this task

You must configure the supported entity types before you can manage this account with Users and Groups in the administrative console. The supported entity types are Group, OrgContainer, and PersonAccount. A Group entity represents a simple collection of entities that might not have any relational context. An OrgContainer entity represents an organization, such as a company or an enterprise, a subsidiary, or an organizational unit, such as a division, a location, or a department. A PersonAccount entity represents a human being. You cannot add or delete the supported entity types, because these types are predefined.

The Base entry for the default parent determines the repository location where entities of the specified type are placed on write operations by user and group management.

Note: To manage users and groups, click **Users and Groups** in the console navigation tree. Click either **Manage Users** or **Manage Groups**.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Click **Supported entity types** to view a list of predefined entity types.
4. Click the name of a predefined entity type to change its configuration.
5. Supply the distinguished name of a base entry in the repository in the Base entry for the default parent field. This entry determines the default location in the repository where entities of this type are placed on write operations by user and group management.
6. Supply the relative distinguished name (RDN™) properties for the specified entity type in the Relative Distinguished Name properties field. Possible values are `cn` for Group, `uid` or `cn` for PersonAccount, and `o`, `ou`, `dc`, and `cn` for OrgContainer. Delimit multiple properties for the OrgContainer entity with a semicolon (;).

The following list outlines known requirements and limitations that apply to specific Lightweight Directory Access Protocol (LDAP) servers:

Using Microsoft Active Directory as the LDAP server

- Unless you modify the LDAP schema to use `uid`, you must specify `cn` in the Relative Distinguished Name (RDN) properties field for the PersonAccount entity type.
- Secure Sockets Layer communications must be enabled to create users with passwords. To select the **Require SSL communications** option, see the topic “Configuring Lightweight Directory Access Protocol in a federated repository configuration” on page 190.
- Typically the value of `user` is specified as the value in the Object classes field for the PersonAccount entity type and the value of `group` is specified as the value in the Object classes field for the Group entity type.

Using a Lotus Domino Enterprise Server as the LDAP server

- Typically, the value of `cn` is specified in the Relative Distinguished Name (RDN) properties field for the PersonAccount entity type. The value of `uid` is also acceptable.

- Typically, both inetOrgPerson and dominoPerson are used as values in the Object classes field for the PersonAccount entity type.

Using Sun ONE Directory Server as the LDAP server

- Typically, groupOfUniqueNames is specified as the value in the Object classes field for the Group entity type.

7. Click **OK**.

Results

After completing these steps, your federated repository configuration, which uses supported entity types, is configured.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** on the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Supported entity types collection:

Use this page to list entity types that are supported by the member repositories or to select an entity type to view or change its configuration properties.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Additional properties, click **Supported entity types**.

You must configure the supported entity types before you can manage this account with Users and Groups in the administrative console. The Base entry for the default parent determines the repository location where entities of the specified type are placed on write operations by user and group management.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Related reference

“Lightweight Directory Access Protocol entity types settings” on page 229
Use this page to configure Lightweight Directory Access Protocol (LDAP) entity types that are supported by the member repositories.

Entity type:

Specifies the entity type name.

Base entry for the default parent:

Specifies the distinguished name of a base entry in the repository.

This entry determines the default location in the repository where entities of this type are placed on write operations by user and group management.

Relative Distinguished Name properties:

Specifies the relative distinguished name (RDN) properties for the specified entity type.

Possible values are `cn` for Group, `uid` or `cn` for PersonAccount, and `o`, `ou`, `dc`, and `cn` for OrgContainer. Delimit multiple properties for the OrgContainer entity with a semicolon (;).

Supported entity types settings:

Use this page to configure entity types that are supported by the member repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Additional properties, click **Supported entity types**.
4. Click the name of a configured entity type to view or change its configuration.

You must configure the supported entity types before you can manage this account with Users and Groups in the administrative console. The Base entry for the default parent determines the repository location where entities of the specified type are placed on write operations by user and group management.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Entity type:

Specifies the name of the entity type.

Base entry for the default parent:

Specifies the distinguished name of a base entry in the repository.

This entry determines the default location in the repository where entities of this type are placed on write operations by user and group management.

Relative Distinguished Name properties:

Specifies the relative distinguished name (RDN) properties for the specified entity type.

Possible values are `cn` for Group, `uid` or `cn` for PersonAccount, and `o`, `ou`, `dc`, and `cn` for OrgContainer. Delimit multiple properties for the OrgContainer entity with a semicolon (;).

Managing repositories in a federated repository configuration

Follow this topic to manage repositories in a federated repository configuration.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.

3. Under Related items, click **Manage repositories**. Repositories that are configured in the system are listed in the collection panel. This list includes repositories that are configured using the federated repository functionality as well as repositories that are created using wsadmin commands described in the topic “IdMgrRepositoryConfig command group for the AdminTask object” on page 1163.
4. Optional: Click **Add** to configure a new external repository. The Lightweight Directory Access Protocol (LDAP) repository configuration settings are described in detail in “Configuring Lightweight Directory Access Protocol in a federated repository configuration” on page 190.

Note: You cannot add a database repository using the administrative console. This repository configuration is supported by using wsadmin commands only.

5. Optional: Click **Delete** to delete a repository that you specified previously using the administrative console or wsadmin commands.

Note: You cannot delete the built-in, file-based repository from the collection panel.

6. Optional: Select one of the LDAP repository identifier entries to view or update an external repository that is configured in the system previously. The steps to configure LDAP settings are described in detail in “Configuring Lightweight Directory Access Protocol in a federated repository configuration” on page 190.

Note: While database repositories that are configured in the system are listed in the collection panel, you cannot update a database repository using the administrative console. Updates to a database repository are supported by using wsadmin commands only.

7. Click **OK**.

Results

After completing these steps, the collection panel under Managing repositories reflects a current list of repositories that are configured in your system.

What to do next

1. To add one or more external repositories that are listed on this collection panel into the realm, see “Managing the realm in a federated repository configuration” on page 170.
2. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
3. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
4. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Note: Be aware of any changes made to built-in, file-based repositories that are part of the federated repositories configuration. You must replicate these changes to the managed nodes to ensure that your network deployment configuration is synchronized. See “Replicating changes to a built-in, file-based repository” for the steps needed to perform this replication.

Replicating changes to a built-in, file-based repository:

Changes to built-in, file-based repositories are not automatically replicated to managed nodes in a federated repositories configuration. You need to use the administrative console to replicate the changes you make to a built-in, file-based repository.

About this task

The network deployment support in a federated repositories configuration only updates the in-memory state of the processes that are running on the managed nodes. Because WebSphere Application Server synchronizes the file systems, the network deployment support does not attempt to update the file systems of the managed nodes.

You must synchronize the node configuration to replicate the changes to the built-in, file-based repository.

1. In the administrative console, click **System Administration > Nodes**. to access the nodes panel.
2. On the Nodes panel, select all the relevant nodes for which the changes to the built-in, file-based repository need to be made.
3. Click **Full Resynchronize**. The resynchronize operation resolves conflicts among configuration files and can take several minutes to complete.

Results

After completing these steps, your federated repository configuration of managed nodes reflects the changes to the built-in, file-based repository.

Manage repositories collection:

Use this page to list repositories that are configured in the system or to select a repository to view or change its configuration properties. You can add or delete external repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Repository identifier:

Specifies a unique identifier for the repository. This identifier uniquely identifies the repository within the cell.

Repository type:

Specifies the repository type, such as File or LDAP.

Repository reference settings:

Use this page to configure a repository reference. A repository reference is a single repository that contains a set of identity entries that are referenced by a base entry into the directory information tree.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Click **Add base entry to realm**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Repository:

Specifies a unique identifier for the repository. This identifier uniquely identifies the repository within the cell.

Expand the drop-down list to display a list of previously defined repository identifiers.

Distinguished name that uniquely identifies this set of entries in the realm:

Specifies the distinguished name (DN) that uniquely identifies this set of entries in the realm.

If multiple repositories are included in the realm, it is necessary to define an additional distinguished name that uniquely identifies this set of entries within the realm. Overlapping base entries are not supported. You should not define two base entries where one is `c=us`, and the other is `o=myorg,c=us` in the same realm; otherwise a search returns duplicate results.

Distinguished name of a base entry in this repository:

Specifies the Lightweight Directory Access Protocol (LDAP) distinguished name (DN) of the base entry within the repository. The entry and its descendents are mapped to the subtree that is identified by the unique base name entry field.

If this field is left blank, then the subtree defaults to the root of the LDAP repository.

Increasing the performance of the federated repository configuration

Follow this page to manage the realm in a federated repository configuration.

Before you begin

The settings that are available on the Performance panel are independent options that pertain specifically to the federated repositories functionality. These options do not affect your entire WebSphere Application Server configuration.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories > repository_name**.
4. Under Additional properties, click **Performance**.
5. Optional: Select the **Limit search time** option and enter the maximum number of milliseconds that the Application Server can use to search through your Lightweight Directory Access Protocol (LDAP) entries.
6. Optional: Select the **Limit search returns** option and enter the maximum number of entries to return that match the search criteria.
7. Optional: Select the **Use connection pooling** option to specify whether the Application Server can store separate connections to the LDAP server for reuse.
8. Optional: Select the **Enable context pool** option to specify whether multiple applications can use the same connection to the LDAP server. If you select the option, specify the initial, preferred, and maximum number of entries that can use the same connection. The **Enable context pool** option can be enabled either in conjunction with the **Use connection pool** option or separately. If this option is

disabled, a new connection is created for each context. You can also select the **Context pool times out** option and specify the number of seconds after which the entries in the context pool expire.

- Optional: Select the **Cache the attributes** option and specify the maximum number of search attribute entries. This option enables WebSphere Application Server to save the LDAP entries so that it can search the entries locally rather than making multiple calls to the LDAP server. Click the **Cache times out** option that is associated with the **Cache the attributes** option to specify the maximum number of seconds that the Application Server can save these entries.
- Optional: Select the **Cache the search results** option and specify the maximum number of search result entries. This option enables WebSphere Application Server to save the results of a search inquiry instead of making multiple calls to the LDAP server to search and retrieve the results of that search. Click the **Cache times out** option that is associated with the **Cache the search results** option to specify the maximum number of seconds that the Application Server can save the results.

Results

These options are available to potentially increase the performance of your federated repositories configuration. However, the any increase in performance is dependant upon your specific configuration.

Lightweight Directory Access Protocol performance settings:

Use this page to minimize impacts to performance by adding opened connections and contexts to internally maintained pools and reusing them. Also minimize performance impacts by maintaining internal caches of retrieved data.

To view this administrative console page, complete the following steps:

- In the administrative console, click **Security > Global security**.
- Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
- Under Related items, click **Manage repositories**.
- Click **Add** to specify a new external repository or select an external repository that is preconfigured.
- Under Additional properties, click **Performance**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Limit search time:

Specifies the timeout value in milliseconds for a Lightweight Directory Access Protocol (LDAP) server to respond before stopping a request.

Data type:	Integer
Units:	Milliseconds
Default:	0
Range:	Equal to or greater than 0. A value of 0 specifies that no search time limit exists.

Limit search returns:

Specifies the maximum number of entries that are returned in a search result.

Data type: Integer
Units: Entries
Default: 0
Range: Equal to or greater than 0. A value of 0 specifies that no search return limit exists.

Use connection pooling:

Specifies whether to utilize the connection pooling function, which is provided in the Software Development Kit (SDK).

Connection pooling is maintained by the Java run time. It is configured by system properties.

Default: Disabled
Range: Enabled or Disabled

Enable context pool:

Specifies whether context pooling is enabled to the LDAP server. To improve performance, use the context pool in combination with connection pooling.

Default: Enabled
Range: Enabled or Disabled

Initial size:

Specifies the number of context instances in the pool when the pool is initially created by the LDAP repository.

Data type: Integer
Default: 1
Range: 1 to 50

Preferred size:

Specifies the preferred number of context instances that the context pool maintains. Both in-use and idle context instances contribute to this number.

Data type: Integer
Default: 3
Range: 0 to 100

Maximum size:

Specifies the maximum number of context instances that can be maintained concurrently by the context pool. Both in-use and idle context instances contribute to this number.

When the pool size reaches the maximum size, no new context instances can be created for a new request. The new request is blocked until a context instance is released or removed. The request periodically checks for context instances that are available in the pool. A request for a pooled context instance uses an existing pooled and idle context instance or a newly created pooled context instance.

A maximum pool size of 0 indicates that the context pool can maintain an infinite number of context instances.

Data type: Integer
Default: 0

Context pool times out:

Specifies the number of seconds for the context pool to time out and remove idle context instances.

A timeout value of 0 indicates that the context pool does not time out context instances.

Data type: Integer
Default: 0

Cache the attributes:

Specifies whether to cache the attributes that are returned from the LDAP server.

Default: Enabled
Range: Enabled or Disabled

Cache size:

Specifies the maximum size of the cache.

Data type: Integer
Default: 4000
Range: Equal to or greater than 100

Cache times out:

Specifies the maximum number of seconds that the cached search results can stay in the cache.

A timeout value of 0 indicates that the cached search results stay in the cache until update operations are made.

Data type: Integer
Units: Seconds
Default: 1200
Range: Equal to or greater than 0

Cache the search results:

Specifies whether to cache the search results that are returned from the LDAP server.

Default: Enabled
Range: Enabled or Disabled

Cache size:

Specifies the maximum size of the cache.

Data type:	Integer
Default:	2000
Range:	Equal to or greater than 100

Cache times out:

Specifies the maximum number of seconds that the cached search results can stay in the cache.

A timeout value of 0 indicates that the cached search results stay in the cache until update operations are made.

Data type:	Integer
Units:	Seconds
Default:	600
Range:	Equal to or greater than 0

Using custom adapters for federated repositories

When the custom adapters for federated repositories are part of the default realm, the users and groups can be managed using wsadmin commands or the administrative console.

About this task

If custom adapters for federated repositories are part of the default realm, you use the administrative console to manage the users and groups in the realm.

Note: The default parent for PersonAccount and Group entities needs to be the same as the base entry of the custom adapter.

To view this administrative console page, complete the following steps:

- In the administrative console, click **Security > Global security**.
- Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
- Under Additional properties, click **Supported entity types**.

You must configure the supported entity types before you can manage this account with Users and Groups in the administrative console. The Base entry for the default parent determines the repository location where entities of the specified type are placed on write operations by user and group management.

1. In the administrative console, click **Users and Groups** to access users and groups panel.
2. Click **Manage Groups** to test the basic functions of the custom adapter with respect to custom adapters for federated repositories.
3. Click **Manage Users** to test the basic functions of the custom adapter with respect to custom adapters for federated repositories.

Results

After completing these steps, you will have ensured that the custom adapter is being used properly.

What to do next

Adjustments to the custom adapter can be made by using the wsadmin tool to make configuration changes. See “Configuring custom adapters for federated repositories using wsadmin” on page 225 for more details.

Sample custom adapters for federated repositories examples:

Out of the box adapters for federated repositories provide File, LDAP, and Database adapters for your use. These adapters implement the `com.ibm.wsspi.wim.Repository` software programming interface (SPI). A virtual member manager custom adapter needs to implement the same SPI.

Developing custom adapters for federated repositories

Out of the box adapters for federated repositories provide File, LDAP and Database adapters for your use. All these adapters implement the `com.ibm.wsspi.wim.Repository` SPI. See the `com.ibm.wsspi.wim.Repository` SPI for more information. As you develop a virtual member manager custom adapter, you need to implement the same SPI.

Custom adapters for federated repositories must not depend on any WebSphere Application Server components, such as data sources and enterprise beans. These WebSphere Application Server components require that security is initialized and enabled prior to startup. If your implementation of custom adapters for federated repositories needs to use data sources to connect to a database, you need to use Java database connectivity (JDBC) to make the connection during server startup. Then, at a later time, switch to using the data sources when the data source is available.

There are examples of suggested behavior and requirements of custom adapters for federated repositories that you can find in the sample code.

A sample custom adapter for federated repositories

A sample custom adapter implementation has been provided as an example. The custom adapter is based on file repository. The sample source code and class files are bundled in `vmmsampleadapter.jar`. The `vmmsampleadapter.jar` can be downloaded at this location: <http://www.ibm.com/developerworks/websphere/downloads/samples/vmmsampleadapter.html>.

Contents of the `vmmsampleadapter.jar` file are as follows:

- Class files for the sample adapter:
 - `com/ibm/ws/wim/adapter/sample/AbstractAdapterImpl.class`
 - `com/ibm/ws/wim/adapter/sample/SampleFileAdapter.class`
 - `com/ibm/ws/wim/adapter/sample/XPathHelper.class`
- Source code for the sample adapter:
 - `src/com/ibm/ws/wim/adapter/sample/AbstractAdapterImpl.java`
 - `src/com/ibm/ws/wim/adapter/sample/SampleFileAdapter.java`
 - `src/com/ibm/ws/wim/adapter/sample/XPathHelper.java`

Note: The sample files should not be used in the production environment. You should make a copy of these files, rename them, and update them based on your specific adapter implementation. Refer to the Javadoc in the source code for more information.

`com/ibm/ws/wim/sample/adapter/AbstractAdapterImpl.java`

Provides an abstract implementation class which handles most of the repository independent internal operations for the adapter and defines some simple abstract methods that should be implemented by the custom adapter. For most cases, you may not need to change this file.

`com/ibm/ws/wim/sample/adapter/SampleFileAdapter.java`

Extends from the `AbstractAdapterImpl` class and implements the abstracts method. This class implements the abstract methods using file as the repository. Adapter providers can use this class as a reference to implement these methods specific to their adapters.

com/ibm/ws/wim/sample/adapter/XPathHelper.java

Defines a helper class to parse the XPath search expression and build the search tree. This helper class also contains the method to evaluate the search expression. If your repository supports a search expression, then you need to convert XPath expression to an expression that your repository can process and let your repository evaluate the expression. This helper class evaluates the search expression based on the use of dataobjects. You can overwrite the `evaluate()` method to perform the evaluation using other objects, such as **java.util.Map**.

Some utility classes have been provided to help adapter providers. Most of these utility methods are used in the sample adapter. Refer to <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp?topic=/com.ibm.websphere.javadoc.doc/vmm/com/ibm/wsspi/wim/package-summary.html> for more details.

Establishing custom adapters for federated repositories

Out of the box adapters for federated repositories provide File, LDAP, and Database adapter for your use. These adapters implement the `com.ibm.wsspi.wim.Repository` software programming interface (SPI). Custom adapters for federated repositories need to implement the same SPI.

Before you begin

Refer to the Repository SPI implementation information in the related references for information about the custom adapters for federated repositories SPI.

Refer to the sample custom adapter code that is available in the `vmmsampleadapter.jar` file. The JAR file contains the sample customer adapter code in the `src/` directory. The `vmmsampleadapter.jar` can be downloaded at this location: <http://www.ibm.com/developerworks/websphere/library/samples/vmmsampleadapter.html>

Note:

- The sample that is provided is intended to familiarize you with the features of custom adapters for federated repositories and the handling of various types of dataobjects. Do not use this sample in an actual production environment.
- Copy the `AbstractAdapterImpl` class and rename it before making changes. Make sure that the new name is appropriate for your adapter.

Custom adapters for federated repositories must not depend on any WebSphere Application Server components, such as data sources and enterprise beans. These WebSphere Application Server components require that security is initialized and enabled prior to startup. If your implementation of the virtual member manager custom adapter needs to use data sources to connect to a database, you need to use Java database connectivity (JDBC) to make the connection during server startup. Then, at a later time, switch to using the data sources when the data source is available.

1. Build your implementation.

Note: EMF JAR files contain version number in their names, such as `v200607270021`. Make sure to change the version number to reflect your installation.

To compile your code, you need the following JAR files in the classpath:

- `app_server_root/plugins/com.ibm.ws.runtime_6.1.0.jar`
- `app_server_root/plugins/org.eclipse.emf.commonj.sdo_2.1.0.v200607270021.jar`
- `app_server_root/plugins/org.eclipse.emf.ecore_2.2.1.v200607270021.jar`
- `app_server_root/plugins/org.eclipse.emf.common_2.2.1.v200607270021.jar`
- `app_server_root/plugins/org.eclipse.emf.ecore.xmi_2.2.0.v200607270021.jar`
- `app_server_root/plugins/org.eclipse.emf.ecore.sdo_2.2.0.v200607270021.jar`

Here is an example:

```

app_server_root/java/bin/javac -classpath
app_server_root/plugins/com.ibm.ws.runtime_6.1.0.jar;
app_server_root/plugins/org.eclipse.emf.commonj.sdo_2.1.0.
v200607270021.jar;app_server_root
/plugins/org.eclipse.emf.ecore_2.2.1.v200607270021.jar;
app_server_root/plugins/org.eclipse.emf.common_2.2.1.
v200607270021.jar;app_server_root/plugins/org.
eclipse.emf.ecore.xmi_2.2.0.v200607270021.jar;
app_server_root/plugins/org.eclipse.emf.ecore.sdo_2.2.0.v200607270021.jar
your_implementation_file.java

```

2. Copy the generated class files or the packaged JAR file to the product classpath. The preferred location is the *app_server_root/lib/ext* directory. This should be copied to the classpaths of all the product processes (cell and all NodeAgents).
3. Configure your custom adapter by following the steps in “Configuring custom adapters for federated repositories using wsadmin.”
4. Test your custom adapter by following the steps in “Using custom adapters for federated repositories” on page 222

What to do next

“Configuring custom adapters for federated repositories using wsadmin” provides details about configuring your custom adapter with the wsadmin tool.

Related tasks

“Configuring custom adapters for federated repositories using wsadmin”

You can use the Jython or Jacl scripting language with the wsadmin tool to define custom adapters in the federated repositories configuration file.

“Using custom adapters for federated repositories” on page 222


When the custom adapters for federated repositories are part of the default realm, the users and groups can be managed using wsadmin commands or the administrative console.

Related reference

“Sample custom adapters for federated repositories examples” on page 223

Out of the box adapters for federated repositories provide File, LDAP, and Database adapters for your use. These adapters implement the `com.ibm.wsspi.wim.Repository` software programming interface (SPI). A virtual member manager custom adapter needs to implement the same SPI.

Related information

 <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp?topic=/com.ibm.websphere.javadoc.doc/vmm/com/ibm/wsspi/wim/Repository.html>

Configuring custom adapters for federated repositories using wsadmin

You can use the Jython or Jacl scripting language with the wsadmin tool to define custom adapters in the federated repositories configuration file.

Before you begin

Shut down the WebSphere Application Server and the wsadmin command window.

About this task

The federated repositories configuration file, `wimconfig.xml`, is shipped with WebSphere Application Server 6.1.x and is located in the *app_server_root/profiles/profile_name/config/cells/cell_name/wim/config* directory.

Note: For additional information about the commands to use for this topic, see “IdMgrRepositoryConfig command group for the AdminTask object” on page 1163.

Use the following steps to add a custom adapter to any federated repositories configuration file and to any realm defined within the configuration file.

1. Open the wimconfig.xml file with a text editor.
2. Add a new config:repositories element to the file. This element should be placed before the config:realmConfiguration element.

The following example configures a custom repository to use the **com.ibm.ws.wim.adapter.sample.SampleFileAdapter** class and sets the **SampleFileRepository** repository as the identifier:

```
<config:repositories adapterClassName="com.ibm.ws.wim.adapter.sample.SampleFileAdapter"
id="SampleFileRepository"/>
```

3. Save the wimconfig.xml file and close the text editor.
4. Copy the vmmsampleadapter.jar file that is provided to app_server_root/lib.
5. Start the wsadmin command prompt application and enter the following command:
wsadmin -conntype none
6. Disable paging in the common repository configuration. Set the supportPaging parameter for the updateIdMgrRepository command to false to disable paging.

Note: You must perform this step because the sample adapter does not support paging.

The following examples use the **SampleFileRepository** repository as the identifier for the custom repository.

Using Jython:

```
AdminTask.updateIdMgrRepository('-id SampleFileRepository -supportPaging false')
```

Using Jacl:

```
$AdminTask updateIdMgrRepository {-id SampleFileRepository -supportPaging
false}
```

Note: A warning will appear until the configuration of the sample repository is complete.

7. Add the necessary custom properties for the adapter. Use the setIdMgrCustomProperty command repeatedly to add multiple properties. Use this command once per property to add multiple properties to your configuration. You must use both the **name** and **value** parameters to add the custom property for the specified repository. For example, to add a custom property of **fileName**, enter the following command.

Using Jython:

```
AdminTask.setIdMgrCustomProperty('-id SampleFileRepository -name fileName
-value "c:\sampleFileRegistry.xml"')
```

Using Jacl:

```
$AdminTask setIdMgrCustomProperty {-id SampleFileRepository -name fileName
-value "c:\sampleFileRegistry.xml"}
```

8. Add a base entry to the adapter configuration. Use the addIdMgrRepositoryBaseEntry command to specify the name of the base entry for the specified repository. For example:

Using Jython:

```
AdminTask.addIdMgrRepositoryBaseEntry('-id SampleFileRepository -name
o=sampleFileRepository')
```

Using Jacl:

```
$AdminTask addIdMgrRepositoryBaseEntry {-id SampleFileRepository -name
o=sampleFileRepository}
```

9. Use the addIdMgrRealmBaseEntry command to add the base entry to the realm, which will link the realm with the repository:

Using Jython:

```
AdminTask.addIdMgrRealmBaseEntry('-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository')
```

Using Jacl:

```
$AdminTask addIdMgrRealmBaseEntry {-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository}
```

10. Save your configuration changes. Enter the following commands to save the new configuration and close the wsadmin tool.

Using Jython:

```
AdminConfig.save()
exit
```

Using Jacl:

```
$AdminConfig save
exit
```

The following example displays the complete text of the newly-revised wimconfig.xml file:

```
<!--
  Begin Copyright

  Licensed Materials - Property of IBM

  virtual member manager

  (C) Copyright IBM Corp. 2005 All Rights Reserved.

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

  End Copyright
-->
<sdo:datagraph xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:config="http://www.ibm.com/websphere/wim
/config" xmlns:sdo="commonj.sdo">
  <config:configurationProvider maxPagingResults="500" maxSearchResults="4500"
maxTotalPagingResults="1000"
pagedCacheTimeOut="900" pagingEntityObject="true" searchTimeOut="600000">
  <config:dynamicModel xsdFileName="wimdatagraph.xsd"/>
  <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="Group">
    <config:rdnProperties>cn</config:rdnProperties>
  </config:supportedEntityTypes>
  <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="OrgContainer">
    <config:rdnProperties>o</config:rdnProperties>
    <config:rdnProperties>ou</config:rdnProperties>
    <config:rdnProperties>dc</config:rdnProperties>
    <config:rdnProperties>cn</config:rdnProperties>
  </config:supportedEntityTypes>
  <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="PersonAccount">
    <config:rdnProperties>uid</config:rdnProperties>
  </config:supportedEntityTypes>
  <config:repositories xsi:type="config:FileRepositoryType" adapterClassName="com.ibm.
ws.wim.adapter.file.was.FileAdapter"
id="InternalFileRepository" supportPaging="false" supportSorting="false" messageDigestAlgorithm="SHA-1">
    <config:baseEntries name="o=defaultWIMFileBasedRealm"/>
  </config:repositories>
  <config:repositories adapterClassName="com.ibm.ws.wim.adapter.sample.SampleFileAdapter"
id="SampleFileRepository">
    <config:CustomProperties name="fileName" value="c:\sampleFileRegistry.xml"/>
    <config:baseEntries name="o=sampleFileRepository"/>
  </config:repositories>
  <config:realmConfiguration defaultRealm="defaultWIMFileBasedRealm">
    <config:realms delimiter="@" name="defaultWIMFileBasedRealm" securityUse="active">
      <config:participatingBaseEntries name="o=defaultWIMFileBasedRealm"/>
      <config:participatingBaseEntries name="o=sampleFileRepository"/>
      <config:uniqueUserIdMapping propertyForInput="uniqueName" propertyForOutput="uniqueName"/>
      <config:userSecurityNameMapping propertyForInput="principalName" propertyForOutput="principalName"/>
      <config:userDisplayNameMapping propertyForInput="principalName" propertyForOutput="principalName"/>
      <config:uniqueGroupIdMapping propertyForInput="uniqueName" propertyForOutput="uniqueName"/>
      <config:groupSecurityNameMapping propertyForInput="cn" propertyForOutput="cn"/>
    </config:realms>
  </config:realmConfiguration>
</sdo:datagraph>
```



```

        <config:groupDisplayNameMapping propertyForInput="cn" propertyForOutput="cn"/>
    </config:realms>
</config:realmConfiguration>
</config:configurationProvider></sdo:datagraph>

```

11. Restart the application server.

Related reference

“Sample custom adapters for federated repositories examples” on page 223

Out of the box adapters for federated repositories provide File, LDAP, and Database adapters for your use. These adapters implement the `com.ibm.wsspi.wim.Repository` software programming interface (SPI). A virtual member manager custom adapter needs to implement the same SPI.

Configuring Lightweight Directory Access Protocol entity types in a federated repository configuration

Follow this task to configure Lightweight Directory Access Protocol (LDAP) entity types in a federated repository configuration.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured. During LDAP configuration, based on the selected LDAP server type, some defaults and mappings are set in the configuration. When the selected LDAP server type is custom, no default is set, and you must set all of the mappings manually. To avoid setting all of the mappings manually, choose a non-custom LDAP server type (for example, IBM Directory Server or SunOne) which matches closely to your LDAP server.

Note: If you click **Add** to specify a new external repository, you must first complete the required fields and click **Apply** before you can proceed to the next step.

5. Under Additional properties, click **LDAP entity types**.
6. View the entity types that are supported by the member repositories, or select an entity type to view or change its configuration properties.
7. Supply the object classes that are mapped to this entity type in the Object classes field. LDAP entries that contain one or more of the object classes belong to this entity type.
8. Supply the search bases that are used to search this entity type. The search bases specified must be subtrees of the base entry in the repository. For example, you can specify the following search bases, where `o=ibm,c=us` is the base entry in the repository:

```
o=ibm,c=us or cn=users,o=ibm,c=us or ou=austin,o=ibm,c=us
```

In the preceding example, you cannot specify search bases `c=us` or `o=ibm,c=uk`.

Delimit multiple search bases with a semicolon (;). For example:

```
ou=austin,o=ibm,c=us;ou=raleigh,o=ibm,c=us
```

9. Supply the LDAP search filter that is used to search this entity type. For example, use `(objectclass=ePerson)` to search for users or `(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames))` to search for groups in an external LDAP repository.

If a search filter is not specified, the object classes and the relative distinguished name (RDN) properties are used to generate the search filter. For information on RDN properties, see “Configuring supported entity types in a federated repository configuration” on page 213.

Results

After completing these steps, LDAP entity types are configured for your LDAP repository.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Lightweight Directory Access Protocol entity types collection:

Use this page to list Lightweight Directory Access Protocol (LDAP) entity types that are supported by the member repositories or to select an LDAP entity type to view or change its configuration properties.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **LDAP entity types**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Entity type:

Specifies the entity type name.

Object classes:

Specifies the object classes that are mapped to this entity type. LDAP entries that contain one or more of the object classes belong to this entity type.

You cannot map multiple entity types to the same LDAP object class.

Lightweight Directory Access Protocol entity types settings:

Use this page to configure Lightweight Directory Access Protocol (LDAP) entity types that are supported by the member repositories.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **LDAP entity types**.
6. Select an entity type to view or change its configuration properties.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Entity type:

Specifies the entity type.

Object classes:

Specifies the object classes that are mapped to this entity type. LDAP entries that contain one or more of the object classes belong to this entity type.

You cannot map multiple entity types to the same LDAP object class.

Search bases:

Specifies the search bases that are used to search this entity type.

The search bases specified must be subtrees of the base entry in the repository. For example, you can specify the following search bases, where `o=ibm,c=us` is the base entry in the repository:

```
o=ibm,c=us or cn=users,o=ibm,c=us or ou=austin,o=ibm,c=us
```

In the preceding example, you cannot specify search bases `c=us` or `o=ibm,c=uk`.

Delimit multiple search bases with a semicolon (;). For example:

```
ou=austin,o=ibm,c=us;ou=raleigh,o=ibm,c=us
```

Search filter:

Specifies the LDAP search filter that is used to search this entity type.

For example, use `(objectclass=ePerson)` to search for users or `(&(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames)))` to search for groups in an external LDAP repository.

If a search filter is not specified, the object classes and the relative distinguished name (RDN) properties are used to generate the search filter.

Configuring group attribute definition settings in a federated repository configuration

Follow this task to configure group definition settings in a federated repository configuration.

Before you begin

Because group attribute definition settings apply only to a Lightweight Directory Access Protocol (LDAP) repository, you must first configure an LDAP repository. For more information, see “Managing repositories in a federated repository configuration” on page 215.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.

Note: If you click **Add** to specify a new external repository, you must first complete the required fields and click **Apply** before you can proceed to the next step.

5. Under Additional properties, click **Group attribute definition**.
6. Supply the name of the group membership attribute in the Name of group membership attribute field. Only one membership attribute can be defined for each LDAP repository.
Every LDAP entry should have this attribute to indicate the groups to which this entry belongs. For example, `memberOf` is the name of the membership attribute that is used in Active Directory. The group membership attribute contains values that reference groups to which this entry belongs. If `UserA` belongs to `GroupA`, then the value of the `memberOf` attribute of `UserA` should contain the distinguished name of `GroupA`.

If your LDAP server does not support the group membership attribute, then do not specify this attribute. The LDAP repository can look up groups by searching the group member attributes, though the performance might be slower.

7. Select the scope of the group membership attribute. The default value is `Direct`.

Direct The membership attribute contains direct groups only. Direct groups are the groups that contain the member. For example, if `Group1` contains `Group2` and `Group2` contains `User1`, then `Group2` is a direct group of `User1`, but `Group1` is not a direct group of `User1`.

Nested

The membership attribute contains both direct groups and nested groups.

All

The membership attribute contains direct groups, nested groups, and dynamic members.

Results

After completing these steps, group attribute definition settings are configured for your LDAP repository.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Group attribute definition settings:

Use this page to specify the name of the group membership attribute. Every Lightweight Directory Access Protocol (LDAP) entry includes this attribute to indicate the group to which this entry belongs.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **Group attribute definition**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Related tasks

“Managing the realm in a federated repository configuration” on page 170
Follow this topic to manage the realm in a federated repository configuration.

Name of group membership attribute:

Specifies the name of the group membership attribute. Only one membership attribute can be defined for each Lightweight Directory Access Protocol (LDAP) repository.

Every LDAP entry should have this attribute to indicate the groups to which this entry belongs. For example, `memberOf` is the name of the membership attribute that is used in Active Directory. The group membership attribute contains values that reference groups to which this entry belongs. If `UserA` belongs to `GroupA`, then the value of the `memberOf` attribute of `UserA` should contain the distinguished name of `GroupA`.

If your LDAP server does not support the group membership attribute, then do not specify this attribute. The LDAP repository can look up groups by searching the group member attributes, though the performance might be slower.

Scope of group membership attribute:

Specifies the scope of the group membership attribute.

Default:	Direct
Range:	Direct The membership attribute contains direct groups only. Direct groups are the groups that contain the member. For example, if <code>Group1</code> contains <code>Group2</code> and <code>Group2</code> contains <code>User1</code> , then <code>Group2</code> is a direct group of <code>User1</code> , but <code>Group1</code> is not a direct group of <code>User1</code> .
	Nested The membership attribute contains both direct groups and nested groups.
	All The membership attribute contains direct groups, nested groups, and dynamic members.

Configuring member attributes in a federated repository configuration

Follow this task to configure member attributes in a federated repository configuration.

Before you begin

Because member attributes apply only to a Lightweight Directory Access Protocol (LDAP) repository, you must first configure an LDAP repository. For more information, see “Managing repositories in a federated repository configuration” on page 215.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.

Note: If you click **Add** to specify a new external repository, you must first complete the required fields and click **Apply** before you can proceed to the next step.

5. Under Additional properties, click **Group attribute definition**.
 6. Under Additional properties, click **Member attributes**.
 7. Click **New** to specify a new member attribute or **Delete** to remove a preconfigured member attribute.
 8. Accept the default, or supply the name of the member attribute in the Name of member attribute field. For example, member and uniqueMember are two commonly used names of member attributes.
- The member attribute is used to store the values that reference members that the group contains. For example, a group type with an object class groupOfNames has a member attribute named member; group type with object class groupOfUniqueNames has a member attribute named uniqueMember. An LDAP repository supports multiple group types if multiple member attributes and their associated group object classes are specified.
9. Supply the object class of the group that uses this member attribute in the Object class field. If this field is not defined, this member attribute applies to all group object classes.
 10. Select the scope of the member attribute. The default value is Direct.

Direct The member attribute contains direct members only. Direct members are members that are directly contained by the group. For example, if Group1 contains Group2 and Group2 contains User1, then User1 is a direct member of Group2, but User1 is not a direct member of Group1.

Nested

The member attribute contains both direct members and nested members.

All

The member attribute contains direct members, nested members, and dynamic members.

Results

After completing these steps, member attributes are configured for your LDAP repository.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Member attributes collection:

Use this page to list Lightweight Directory Access Protocol (LDAP) member attributes or to select a member attribute to view or change its configuration properties.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **Group attribute definition**.
6. Under Additional properties, click **Member attributes**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Name:

Specifies the name of the member attribute in LDAP. For example, member and uniqueMember are two commonly used names of member attributes.

The member attribute is used to store the values that reference members that the group contains. For example, a group type with an object class groupOfNames has a member attribute named member; group type with object class groupOfUniqueNames has a member attribute named uniqueMember. An LDAP repository supports multiple group types if multiple member attributes and their associated group object classes are specified.

Scope:

Specifies the scope of the member attribute.

Default:	Direct
Range:	Direct The member attribute contains direct members only. Direct members are members that are directly contained by the group. For example, if Group1 contains Group2 and Group2 contains User1, then User1 is a direct member of Group2, but User1 is not a direct member of Group1.
	Nested The member attribute contains both direct members and nested members.
	All The member attribute contains direct members, nested members, and dynamic members.

Object class:

Specifies the object class of the group that uses this member attribute. If this field is not defined, this member attribute applies to all group object classes.

Member attributes settings:

Use this page to configure Lightweight Directory Access Protocol (LDAP) member attributes.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is pre-configured.
5. Under Additional properties, click **Group attribute definition**.
6. Under Additional properties, click **Member attributes**.
7. Click **New** to specify a new member attribute.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Name of member attribute:

Specifies the name of the member attribute in LDAP. For example, member and uniqueMember are two commonly used names of member attributes.

The member attribute is used to store the values that reference members that the group contains. For example, a group type with an object class groupOfNames has a member attribute named member; group type with object class groupOfUniqueNames has a member attribute named uniqueMember. An LDAP repository supports multiple group types if multiple member attributes and their associated group object classes are specified.

Object class:

Specifies the object class of the group that uses this member attribute. If this field is not defined, this member attribute applies to all group object classes.

Scope:

Specifies the scope of the member attribute.

Default:	Direct
Range:	Direct The member attribute contains direct members only. Direct members are members that are directly contained by the group. For example, if Group1 contains Group2 and Group2 contains User1, then User1 is a direct member of Group2, but User1 is not a direct member of Group1.
	Nested The member attribute contains both direct members and nested members.
	All The member attribute contains direct members, nested members, and dynamic members.

Configuring dynamic member attributes in a federated repository configuration

Follow this task to configure dynamic member attributes in a federated repository configuration.

Before you begin

Because dynamic member attributes apply only to a Lightweight Directory Access Protocol (LDAP) repository, you must first configure an LDAP repository. For more information, see “Managing repositories in a federated repository configuration” on page 215.

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.

Note: If you click **Add** to specify a new external repository, you must first complete the required fields and click **Apply** before you can proceed to the next step.

5. Under Additional properties, click **Group attribute definition**.
6. Under Additional properties, click **Dynamic member attributes**.
7. Click **New** to specify a new dynamic member attribute or **Delete** to remove a preconfigured dynamic member attribute.

8. Accept the default, or supply the name of the dynamic member attribute in the Name of dynamic member attribute field. The name of the dynamic member attribute defines the filter for dynamic group members in LDAP, for example, memberURL is the name of a commonly used dynamic member attribute.

If both member and dynamic member attributes are specified for the same group type, this group type is a hybrid group with both static and dynamic members.

A dynamic group defines its members differently than a static group. Instead of listing the members individually, the dynamic group defines its members using an LDAP search. The filter for the search is defined in a dynamic member attribute. For example, the dynamic group uses the structural objectclass groupOfURLs, or auxiliary objectclass ibm-dynamicGroup, and the attribute memberURL, to define the search using a simplified LDAP URL syntax:

```
ldap:///<base DN of search> ? ? <scope of search> ? <searchfilter>
```

The following is an example of the LDAP URL that defines all entries that are under o=Acme with the objectclass=person:

```
ldap:///o=Acme,c=US??sub?objectclass=person
```

9. Supply the object class of the group that contains the dynamic member attribute in the Dynamic object class field, for example, groupOfURLs. If this property is not defined, the dynamic member attribute applies to all group object classes.

Results

After completing these steps, dynamic member attributes are configured for your LDAP repository.

What to do next

1. After configuring the federated repositories, click **Security > Global security** to return to the Global security panel. Verify that Federated repositories is identified in the Current realm definition field. If Federated repositories is not identified, select **Federated repositories** from the Available realm definitions field and click **Set as current**. To verify the federated repositories configuration, click **Apply** on the Global security panel. If Federated repositories is not identified in the Current realm definition field, your federated repositories configuration is not used by WebSphere Application Server.
2. If you are enabling security, complete the remaining steps as specified in “Enabling security for the realm” on page 53. As the final step, validate this setup by clicking **Apply** in the Global security panel.
3. Save, stop, and restart all the product servers (deployment managers, nodes, and Application Servers) for changes in this panel to take effect. If the server comes up without any problems, the setup is correct.

Dynamic member attributes collection:

Use this page to manage Lightweight Directory Access Protocol (LDAP) dynamic member attributes.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **Group attribute definition**.
6. Under Additional properties, click **Dynamic member attributes**.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Name:

Specifies the name of the attribute that defines the filter for dynamic group members in LDAP. For example, memberURL is the name of a commonly used dynamic member attribute.

If both member and dynamic member attributes are specified for the same group type, this group type is a hybrid group with both static and dynamic members.

A dynamic group defines its members differently than a static group. Instead of listing the members individually, the dynamic group defines its members using an LDAP search. The filter for the search is defined in a dynamic member attribute. For example, the dynamic group uses the structural objectclass groupOfURLs, or auxiliary objectclass ibm-dynamicGroup, and the attribute memberURL, to define the search using a simplified LDAP URL syntax:

```
ldap:///<base DN of search>??<scope of search>?<searchfilter>
```

The following is an example of the LDAP URL that defines all entries that are under o=Acme with the objectclass=person:

```
ldap:///o=Acme,c=US??sub?objectclass=person
```

Object class:

Specifies the object class of the group that contains this dynamic member attribute, for example, groupOfURLs. If this property is not defined, the dynamic member attribute applies to all group object classes.

Dynamic member attributes settings:

Use this page to configure Lightweight Directory Access Protocol (LDAP) dynamic member attributes.

To view this administrative console page, complete the following steps:

1. In the administrative console, click **Security > Global security**.
2. Under User account repository, select **Federated repositories** from the Available realm definitions field and click **Configure**.
3. Under Related items, click **Manage repositories**.
4. Click **Add** to specify a new external repository or select an external repository that is preconfigured.
5. Under Additional properties, click **Group attribute definition**.
6. Under Additional properties, click **Dynamic member attributes**.
7. Click **New** to specify a new dynamic member attribute.

When you finish adding or updating your federated repository configuration, go to the **Security > Global security** panel and click **Apply** to validate the changes.

Name of dynamic member attribute:

Specifies the name of the attribute that defines the filter for dynamic group members in LDAP. For example, memberURL is the name of a commonly used dynamic member attribute.

If both member and dynamic member attributes are specified for the same group type, this group type is a hybrid group with both static and dynamic members.

A dynamic group defines its members differently than a static group. Instead of listing the members individually, the dynamic group defines its members using an LDAP search. The filter for the search is

defined in a dynamic member attribute. For example, the dynamic group uses the structural objectclass `groupOfURLs`, or auxiliary objectclass `ibm-dynamicGroup`, and the attribute `memberURL`, to define the search using a simplified LDAP URL syntax:

```
ldap:///<base DN of search>??<scope of search>?<searchfilter>
```

The following is an example of the LDAP URL that defines all entries that are under `o=Acme` with the `objectclass=person`:

```
ldap:///o=Acme,c=US??sub?objectclass=person
```

Dynamic object class:

Specifies the object class of the group that contains this dynamic member attribute, for example, `groupOfURLs`. If this property is not defined, the dynamic member attribute applies to all group object classes.

Local operating system registries

With the registry implementation for the local operating system, the WebSphere Application Server authentication mechanism can use the user accounts database of the local operating system.

Lightweight Directory Access Protocol (LDAP) is a centralized registry. Most local operating system registries are not centralized registries.

WebSphere Application Server provides implementations for the Windows local accounts registry and domain registry, as well as implementations for the Linux, Solaris, and AIX user accounts registries. Windows Active Directory is supported through the LDAP user registry implementation discussed later.

Note: For an Active Directory (domain controller), the three group scopes are Domain Local Group, Global Group, and Universal Group. For an Active Directory (Domain Controller), the two group types are Security and Distribution.

When a group is created, the default value is Global and the default type is Security. With Windows NT[®] domain registry support for Windows 2000 and 2003 domain controllers, WebSphere Application Server only supports Global groups that are the Security type. It is recommended that you use the Active Directory registry support rather than a Windows NT domain registry if you use Windows 2000 and 2003 domain controllers because the Active Directory supports all group scopes and types. The Active Directory also supports a nested group that is not supported by Windows NT domain registry. The Active Directory is a centralized control registry.

Note: WebSphere Application Server does not have to install the member of the domain because it can be installed on any machine on any platform. Note that the Windows NT domain native call returns the support group only without an error.

Do not use a local operating system registry in a WebSphere Application Server environment where application servers are dispersed across more than one machine because each machine has its own user registry.

The Windows domain registry and Network Information Services (NIS) are exceptions. Both the Windows domain registry and Network Information Services (NIS) are centralized registries. The Windows domain registry is supported by WebSphere Application Server; however, NIS is not supported.

As mentioned previously, the access IDs taken from the user registry are used during authorization checks. Because these IDs are typically unique identifiers, they vary from machine to machine, even if the exact users and passwords exist on each machine.

Web client certificate authentication is not currently supported when using the local operating system user registry. However, Java client certificate authentication does function with a local operating user registry. Java client certificate authentication maps the first attribute of the certificate domain name to the user ID in the user registry.

Even though Java client certificates function correctly, the following error displays in the SystemOut.log file:

```
CWSCJ0337E: The mapCertificate method is not supported
```

The error is intended for Web client certificates; however, it also displays for Java client certificates. Ignore this error for Java client certificates.

Required privileges

The user that is running the WebSphere Application Server process requires enough operating system privilege to call the Windows systems application programming interface (API) for authenticating and obtaining user and group information from the Windows operating system. This user logs into the machine, or if running as a service, is the Log On As user. Depending on the machine and whether the machine is a standalone machine or a machine that is part of a domain or is the domain controller, the access requirements vary.

- For a standalone machine, the user:
 - Is a member of the administrative group.
 - Has the Act as part of the operating system privilege.
 - Has the Log on as a service privilege, if the server is run as a service.
- For a machine that is a member of a domain, only a domain user can start the server process and:
 - Is a member of the domain administrative groups in the domain controller.
 - Has the Act as part of the operating system privilege in the Domain security policy on the domain controller.
 - Has the Act as part of the operating system privilege in the Local security policy on the local machine.
 - Has the Log on as a service privilege on the local machine, if the server is run as a service.

The user is a domain user and not a local user, which implies that when a machine is part of a domain, only a domain user can start the server.
- For a domain controller machine, the user:
 - Is a member of the domain administrative groups in the domain controller.
 - Has the Act as part of the operating system privilege in the Domain security policy on the domain controller.
 - Has the Log on as a service privilege on the domain controller, if the server is run as a service.

If the user running the server does not have the required privilege, you might see one of the following exception messages in the log files:

- A required privilege is not held by the client.
- Access is denied.

Note: On the Windows Vista operating system or Windows Server 2008, the application server must be started with Administrator privileges if you are using a Command Prompt. Start the application server from a Command Prompt window that is launched by performing the following actions:

- Right-click a Command Prompt shortcut.
- Click **Run As Administrator**.

When you open the Command Prompt window as Administrator, an operating-system dialog appears that asks you if you want to continue. Click **Continue** to proceed.

Domain and local user registries

When WebSphere Application Server is started, the security run-time initialization process dynamically attempts to determine if the local machine is a member of a Windows domain. If the machine is part of a domain then by default both the local registry users or groups and the domain registry users or groups can be used for authentication and authorization purposes with the domain registry taking precedence. The list of users and groups that is presented during the security role mapping includes users and groups from both the local user registry and the domain user registry. The users and groups can be distinguished by the associated host names.

WebSphere Application Server does not support trusted domains.

If the machine is not a member of a Windows system domain, the user registry local to that machine is used.

Using both the domain user registry and the local operating system registry

When the machine that hosts the WebSphere Application Server process is a member of a domain, both the local and the domain user registries are used by default. The following section describes more on this topic and recommends some best practices to avoid unfavorable consequences.

Note: Although this section does not directly describe z/OS considerations, you should be aware that overall security operations are affected by how well you set up these registries.

- **Best practices**

In general, if the local and the domain registries do not contain common users or groups, it is simpler to administer and it eliminates unfavorable side effects. If possible, give users and groups access to unique security roles, including the server ID and administrative roles. In this situation, select the users and groups from either the local user registry or the domain user registry to map to the roles.

In cases where the same users or groups exist in both the local user registry and the domain user registry, it is recommended that at least the server ID and the users and groups that are mapped to the administrative roles be unique in the registries and exist only in the domain.

If a common set of users exists, set a different password to make sure that the appropriate user is authenticated.

- **How it works**

When a machine is part of a domain, the domain user registry takes precedence over the local user registry. For example, when a user logs into the system, the domain user registry tries to authenticate the user first. If authentication fails, the local user registry is used. When a user or a group is mapped to a role, the user and group information is first obtained from the domain user registry. In case of failure, the local user registry is tried.

However, when a fully qualified user or a group name, one with an attached domain or host name, is mapped to a role, only that user registry is used to get the information. Use the administrative console or scripts to get the fully qualified user and group names, which is the recommended way to map users and groups to roles.

Note: A user, Bob, on one machine in the local OS user registry, for example, is not the same as the user Bob on another machine in the domain user registry, for example, because the unique ID of Bob, which is the security identifier [SID] in this case, is different in different user registries.

- **Examples**

The MyMachine machine is part of the MyDomain domain. The MyMachine machine contains the following users and groups:

- MyMachine\user2
- MyMachine\user3
- MyMachine\group2

The MyDomain domain contains the following users and groups:

- MyDomain\user1
- MyDomain\user2
- MyDomain\group1
- MyDomain\group2

Here are some scenarios that assume the previous set of users and groups:

1. When user2 logs into the system, the domain user registry is used for authentication. If the authentication fails because the password is different, for example, the local user registry is used.
2. If the MyMachine\user2 user is mapped to a role, only the user2 user in MyMachine machine has access. Thus, if the user2 password is the same on both the local and the domain user registries, the user2 user cannot access the resource because the user2 user is always authenticated using the domain user registry. If both user registries have common users, it is recommended that you have different passwords.
3. If the group2 group is mapped to a role, only the users who are members of the MyDomain\group2 group can access the resource because group2 information is first obtained from the domain user registry.
4. If the MyMachine\group2 group is mapped to a role, only the users who are members of the MyMachine\group2 group can access the resource. A specific group is mapped to the role (MyMachine\group2 instead of just group2).
5. Use either the user3 user or the MyMachine\user3 user to map to a role because the user3 user is unique as it exists in one user registry only.

Authorizing with the domain user registry first can cause problems if a user exists in both the domain and local user registries with the same password. Role-based authorization can fail in this situation because the user is first authenticated within the domain user registry. This authentication produces a unique domain security ID that is used in WebSphere Application Server during the authorization check. However, the local user registry is used for role assignment. The domain security ID does not match the unique security ID that is associated with the role. To avoid this problem, map security roles to domain users instead of local users.

Only a centralized repository can be used if more than one node is involved. This usage implies that only the domain user registry can be used because the user and group unique IDs (SIDs) differ on various nodes, as previously mentioned.

Using either the local or the domain user registry

. If you want to access users and groups from either the local or the domain user registry, instead of both, set the `com.ibm.websphere.registry.UseRegistry` property. This property can be set to either `local` or `domain`. When this property is set to `local` (case insensitive) only the local user registry is used. When this property is set to `domain`, (case insensitive) only the domain user registry is used.

Set this property by completing the following steps to access the **Custom Properties** panel in the administrative console:

1. Click **Security > Global security**
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Local operating system**, and click **Configure**.
3. Under Additional properties, click **Custom properties**.

You can also use `wsadmin` to configure this property. When the property is set, the privilege requirement for the user who is running the product process does not change. For example, if this property is set to `local`, the user that is running the process requires the same privilege, as if the property was not set.

UNIX

Using UNIX system user registries

When using UNIX system user registries, the process ID that runs the WebSphere Application Server process needs the root authority to call the local operating system APIs for authentication and for obtaining user or group information.

Note: In UNIX systems, only the local machine user registry is used. Network Information Service (NIS) (Yellow Pages) is not supported.

HP-UX

Note: If you are using the local operating system user registry, HP-UX must be configured in untrusted mode. Trusted mode is not supported if using the local operating system user registry.

Linux

Solaris

Using Linux and Solaris system user registries

For WebSphere Application Server local operating system registry to work on the Linux and Solaris platforms, a shadow password file must exist. The shadow password file is named `shadow` and is located in the `/etc` directory. If the shadow password file does not exist, an error occurs after enabling administrative security and configuring the registry as local operating system.

To create the shadow file, run the `pwconv` command (with no parameters). This command creates an `/etc/shadow` file from the `/etc/passwd` file. After creating the shadow file, you can enable local operating system security successfully.

Standalone Lightweight Directory Access Protocol registries

A Standalone Lightweight Directory Access Protocol (LDAP) registry performs authentication using an LDAP binding.

WebSphere Application Server security provides and supports the implementation of most major LDAP directory servers, which can act as the repository for user and group information. These LDAP servers are called by the product processes for authenticating a user and other security-related tasks. For example, the servers are used to retrieve user or group information. This support is provided by using different user and group filters to obtain the user and group information. These filters have default values that you can modify to fit your needs. The custom LDAP feature enables you to use any other LDAP server, which is not in the product-supported list of LDAP servers, for its user registry by using the appropriate filters.

To use LDAP as the user registry, you need to know an administrative user name that is defined in the registry, the server host and port, the base distinguished name (DN) and, if necessary, the bind DN and the bind password. You can choose any valid user in the registry that is searchable and have administrative privileges. In some LDAP servers, the administrative users are not searchable and cannot be used, for example, `cn=root` in SecureWay. This user is referred to as WebSphere Application Server security server ID, server ID, or server user ID in the documentation. Being a server ID means a user has special privileges when calling some protected internal methods. Normally, this ID and password are used to log into the administrative console after security is turned on. You can use other users to log in if those users are part of the administrative roles.

When security is enabled in the product, the primary administrative user name and password are authenticated with the registry during the product startup. If authentication fails, the server does not start. It is important to choose an ID and password that do not expire or change often. If the product server user ID or password need to change in the registry, make sure that the changes are performed when all the product servers are up and running.

When the changes are done in the registry, use the steps that are described in “Configuring Lightweight Directory Access Protocol user registries” on page 119. Change the ID, password, and other configuration information, save, stop, and restart all the servers so that the new ID or password is used by the product. If any problems occur starting the product when security is enabled, disable security before the server can start up. To avoid these problems, make sure that any changes in this panel are validated in the Global security panel. When the server is up, you can change the ID, password, and other configuration information and then enable security.

You can use the custom Lightweight Directory Access Protocol (LDAP) feature to support any LDAP server by setting up the correct configuration. However, support is not extended to these custom LDAP servers because many configuration possibilities exist.

The users and groups and security role mapping information is used by the configured authorization engine to perform access control decisions.

Dynamic groups and nested group support

Dynamic and nested groups simplify WebSphere Application Server security management and increase its effectiveness and flexibility.

Dynamic groups contain a group name and membership criteria:

- The group membership information is as current as the information on the user object.
- There is no need to manually maintain members on the group object.
- Dynamic groups are designed so an application does not need a large amount of information from the directory to find out if someone is a member of a group.

Nested groups enable the creation of hierarchical relationships that are used to define inherited group membership. A nested group is defined as a child group entry whose distinguished name (DN) is referenced by a parent group entry attribute.

You only need to assign a larger parent group if all nested groups share the same privilege. Assigning a role to a single parent group simplifies the run-time authorization table.

Dynamic groups and nested group support for the IBM Tivoli Directory Server

WebSphere Application Server supports all Lightweight Directory Access Protocol (LDAP) dynamic and nested groups when using IBM Tivoli Directory Server. This function is enabled by default by taking advantage of a new feature in IBM Tivoli Directory Server. IBM Tivoli Directory Server uses the `ibm-allGroups` forward-reference group attribute that automatically calculates all the group memberships including dynamic and recursive memberships for a user. Security directly locates a user group membership from a user object rather than indirectly search all the groups to match group members.

For more information, see “Configuring dynamic and nested group support for the IBM Tivoli Directory Server” on page 137.

Dynamic and nested group support for the SunONE or iPlanet Directory Server

The SunONE or iPlanet Directory Server uses two grouping mechanisms:

Groups

Entries that name other entries as a list of members or as a filter for members.

Roles Entries that name other entries as a list of members or as a filter for members. Additional functionality is provided by generating the `nsrole` attribute on each role member.

Three types of roles are available:

Filtered roles

Depends upon the attributes that are contained in each entry. Entries are members, if they match a specified Lightweight Directory Access Protocol (LDAP) filter. This role is equivalent to a dynamic group.

Nested roles

Creates roles that contain other roles. This role is equivalent to a nested group.

Managed roles

Explicitly assigns a role to member entries. This role is equivalent to a static group.

Refer to “Configuring dynamic and nested group support for the SunONE or iPlanet Directory Server” on page 137 for more information.

Security failover among multiple LDAP servers

WebSphere Application Server security can be configured to attempt failovers between multiple Lightweight Directory Access Protocol (LDAP) hosts.

If the current active LDAP server is unavailable, WebSphere Application Server security attempts a failover to the first available LDAP host in the specified host list. The multiple LDAP servers can be replicas of the same master LDAP server, or they can be any LDAP host with the same schema, which contain data that is imported from the same LDAP Data Interchange Format (LDIF) file.

Whenever a failover occurs, WebSphere Application Server security always uses the first available LDAP server in the specified host list. For example, if there are four LDAP servers configured in the order of L1, L2, L3, and L4, L1 is treated as the primary LDAP server. The preference of connection is from L1 to L4. If, for example, WebSphere Application Server security is currently connected to L4, and failover or reconnection is necessary, WebSphere Application Server security first attempts to connect to L1, L2, and then L3 in that order until the connection is successful.

The current LDAP host name is logged in message CWSCJ0419I in the WebSphere Application Server log file, `SystemOut.log`. If you want to reconnect to the primary LDAP host, run the WebSphere Application Server MBean method, `resetLDAPBindInfo`, with `null,null` as the input.

To configure LDAP failover among multiple LDAP hosts, you must use `wsadmin` or `ConfigService` to include the backup LDAP host, which does not have a number limitation. The LDAP host that is displayed in the administrative console is the primary LDAP host, and is the first item listed in the LDAP host list in `security.xml`.

The WebSphere Application Server security realm name defaults to the primary LDAP host name that is displayed in the administrative console. It includes a trailing colon and a port number (if one exists). However, the custom property, `com.ibm.websphere.security.ldap.logicRealm`, can be added to override the default security realm name. Use the `logicRealm` name to configure each cell to have its own LDAP host for interoperability and backward compatibility, and to provide flexibility for adding or removing the LDAP host dynamically. If migrating from a previous installation, the new `logicRealm` name does not take effect until administrative security is enabled again. To be compatible with a previous release that does not support logic realm, the `logicRealm` name has to be the same as that used by the previous installation (the LDAP host name, including a trailing colon and port number).

The following Jacl example shows how to use `wsadmin` to add a backup LDAP host for failover:

```
proc LDAPAdd {args} {
  global AdminConfig AdminControl ldapServer ldapPort
  set ldapServer [lindex $args 0]
  set ldapPort [lindex $args 1]
  global ldapUserRegistryId
  if {[catch {$AdminConfig list LDAPUserRegistry} result]} {
    puts stdout "\$AdminConfig list LDAPUserRegistry caught an exception $result\n"
    return
  } else {
    if {$result != {}} {
      set ldapUserRegistryId [lindex $result 0]
    } else {
      return;
    }
  }
  set secMbean [$AdminControl queryNames type=SecurityAdmin,*]
  set Attrs2 [list [list hosts [list [list [list host $ldapServer]
```

```

        [list port $ldapPort]]]]]
$AdminConfig modify $ldapUserRegistryId $Attrs2
$AdminConfig save
}

```

The following Jython example shows how to use wsadmin to add a backup LDAP host for failover:

```

#-----
# Add ldap hostname and port
#   wsadmin -f LDAPAdd.py arg1 arg2
#
# The script expects some parameters:
#   arg1 - LDAP Server hostname
#   arg2 - LDAP Server portnumber
#-----
import java

#-----
# get the line separator and use to do the parsing
# since the line separator on different platform are different
lineSeparator = java.lang.System.getProperty('line.separator')

#-----
# add LDAP host
#-----
def LDAPAdd (ldapServer, ldapPort):
    global AdminConfig, lineSeparator, ldapUserRegistryId
    try:
        ldapObject = AdminConfig.list("LDAPUserRegistry")
        if len(ldapObject) == 0:
            print "LDAPUserRegistry ConfigId was not found\n"
            return

        ldapUserRegistryId = ldapObject.split(lineSeparator)[0]
        print "Got LDAPUserRegistry ConfigId is " + ldapUserRegistryId + "\n"
    except:
        print "AdminConfig.list('LDAPUserRegistry') caught an exception\n"

    try:
        secMbeans = AdminControl.queryNames('WebSphere:type=SecurityAdmin,*')
        if len(secMbeans) == 0:
            print "Security Mbean was not found\n"
            return

        secMbean = secMbeans.split(lineSeparator)[0]
        print "Got Security Mbean is " + secMbean + "\n"
    except:
        print "AdminControl.queryNames('WebSphere:type=SecurityAdmin,*') caught an exception\n"

    attrs2 = [["hosts", [[["host", ldapServer], ["port", ldapPort]]]]]
    try:
        AdminConfig.modify(ldapUserRegistryId, attrs2)
        try:
            AdminConfig.save()
            print "Done setting up attributes values for LDAP User Registry"
            print "Updated was saved successfully\n"
        except:
            print "AdminConfig.save() caught an exception\n"
    except:
        print "AdminConfig.modify(" + ldapUserRegistryId + ", " + attrs2 + ") caught an exception\n"
    return

#-----
# Main entry point
#-----
if len(sys.argv) < 2 or len(sys.argv) > 3:

```

```

        print("LDAPAdd: this script requires 2 parameters: LDAP server hostname and LDAP server port number\n")
        print("e.g.: LDAPAdd ldaphost 389\n")
        sys.exit(1)
else:
    ldapServer = sys.argv[0]
    ldapPort = sys.argv[1]
    LDAPAdd(ldapServer, ldapPort)
    sys.exit(0)

```

Federated repositories

Federated repositories enable you to use multiple repositories with WebSphere Application Server. These repositories, which can be file-based repositories, LDAP repositories, or a sub-tree of an LDAP repository, are defined and theoretically combined under a single realm. All of the user repositories that are configured under the federated repository functionality are invisible to WebSphere Application Server.

When you use the federated repositories functionality, all of the configured repositories, which you specify as part of the federated repository configuration, become active. It is required that the user ID, and the distinguished name (DN) for an LDAP repository, be unique in multiple user repositories that are configured under the same federated repository configuration. For example, there might be three different repositories that are configured for the federated repositories configuration: Repository A, Repository B, and Repository C. When user1 logs in, the federated repository adapter searches each of the repositories for all of the occurrences of that user. If multiple instances of that user are found in the combined repositories, an error message displays.

In addition, the federated repositories functionality in WebSphere Application Server supports the logical joining of entries across multiple user repositories when the Application Server searches and retrieves entries from the repositories. For example, when an application calls for a sorted list of people whose age is greater than twenty, WebSphere Application searches all of the repositories in the federated repositories configuration. The results are combined and sorted before the Application Server returns the results to the application.

Unlike the local operating system, standalone LDAP registry, or custom registry options, federated repositories provide user and group management with read and write capabilities. When you configure federated repositories, you can use one of the following methods to add, create, and delete users and groups:

Note: If you configure multiple repositories under the federated repositories realm, you must also configure supported entity types and specify a base entry for the default parent. The base entry for the default parent determines the repository location where entities of the specified type are placed on write operations by user and group management. See “Configuring supported entity types in a federated repository configuration” on page 213 for details.

- Use the user management application programming interfaces (API). For more information, refer to articles under “Developing with virtual member manager” in this information center.
- Use the administrative console. To manage users and groups within the administrative console, click **Users and Groups > Manage Users** or **Users and Groups > Manage Groups**. For information on user and group management, click the Help link that displays in the upper right corner of the window. From the left navigation pane, click **Users and Groups**.
- Use the wsadmin commands. For more information, see “WIMManagementCommands command group for the AdminTask object” on page 1218.

If you do not configure the federated repositories functionality or do not enable federated repositories as the active repository, you cannot use the user management capabilities that are associated with federated repositories. You can configure an LDAP server as the active user registry and configure the same LDAP server under federated repositories, but not select federated repositories as the active user repository. With this scenario, authentication takes place using the LDAP server, and you can use the user management functionality for the LDAP server that is available for federated repositories.

The following table compares the federated repository functionality that is available in WebSphere Application Server Version 7.0 with the registry functionality that remains unchanged from previous versions of the Application Server.

Table 3. Federated repositories versus user registry implementations

Federated repositories	User registry
Supports multiple types of repositories such as file-based, LDAP, database, and custom. In WebSphere Application Server Version 7.0, file-based and LDAP repositories are supported by the administrative console. However, the federated repositories functionality does not support local operating system implementations. For database and custom repositories, you can use the wsadmin command-line interface or the configuration application programming interfaces (API).	Supports multiple types of registries such as the local operating system, a standalone LDAP registry, and a standalone custom registry.
Supports multiple repositories in a realm within a cell.	Supports one registry only in a realm within a cell.
Provides read and write capabilities for the repositories that are defined in the federated repository configuration.	Provides read only capability for the registries.
Provides account and password policy support as defined by the registry type. However, this support is not provided by the federated repository functionality.	Provides account and password policy support as defined by the registry type.
Supports identity profiles.	Does not support identity profiles.
Uses the custom UserRegistry implementation.	Uses the custom UserRegistry implementation.

Selecting an authentication mechanism

An *authentication mechanism* defines rules about security information, such as whether a credential is forwardable to another Java process, and the format of how security information is stored in both credentials and tokens. You can select and configure an authentication mechanism by using the administrative console.

About this task

Authentication is the process of establishing whether a client is who or what it claims to be in a particular context. A client can be either an end user, a machine, or an application. An authentication mechanism in WebSphere Application Server typically collaborates closely with a *user registry*. The user registry is the user and groups account repository that the authentication mechanism consults with when performing authentication. The authentication mechanism is responsible for creating a *credential*, which is an internal product representation of a successfully authenticated client user. Not all credentials are created equally. The abilities of the credential are determined by the configured authentication mechanism.

WebSphere Application Server provides three authentication mechanisms: Lightweight Third Party Authentication (LTPA), Kerberos, and RSA token authentication mechanism.

Security support for Kerberos as the authentication mechanism has been added for this release of WebSphere Application Server. Kerberos (KRB5) is a mature, flexible, open, and very secure network authentication protocol. Kerberos includes authentication, mutual authentication, message integrity and confidentiality and delegation features. KRB5 is used for Kerberos in the administrative console and in the **sas.client.props**, **soap.client.props** and **ipc.client.props** files.

The RSA token authentication mechanism is new to this release of WebSphere Application Server. It aids the flexible management objective to preserve the base profiles configurations and isolate them from a

security perspective. This mechanism permits the base profiles managed by an administrative agent to have different Lightweight Third-Party Authentication (LTPA) keys, different user registries, and different administrative users.

Note: Simple WebSphere Authentication Mechanism (SWAM) is deprecated in this release. SWAM does not provide authenticated communication between different servers.

Authentication is required for enterprise bean clients and Web clients when they access protected resources. Enterprise bean clients, like a servlet or other enterprise beans or a pure client, send the authentication information to a Web application server using one of the following protocols:

- Common Secure Interoperability Version 2 (CSIv2)
- Secure Authentication Service (SAS)

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Web clients use the HTTP or HTTPS protocol to send the authentication information.

The authentication information can be basic authentication (user ID and password), a credential token, or a client certificate. The Web authentication is performed by the Web Authentication module.

You can configure Web authentication for a Web client by using the administrative console. Click **Security > Global security**. Under Authentication, expand **Web and SIP security** and click **General settings**. The following options exist for Web authentication:

Authenticate only when the URI is protected

Specifies that the Web client can retrieve an authenticated identity only when it accesses a protected Uniform Resource Identifier (URI). WebSphere Application Server challenges the Web client to provide authentication data when the Web client accesses a URI that is protected by a J2EE role. This default option is also available in previous versions of WebSphere Application Server.

Use available authentication data when an unprotected URI is accessed

Specifies that the Web client is authorized to call the `getRemoteUser`, `isUserInRole`, and `getUserPrincipal` methods; retrieves an authenticated identity from either a protected or an unprotected URI. Although the authentication data is not used when you access an unprotected URI, the authentication data is retained for future use. This option is available when you select the **Authenticate only when the URI is protected** check box.

Authenticate when any URI is accessed

Specifies that the Web client must provide authentication data regardless of whether the URI is protected.

Default to basic authentication when certificate authentication for the HTTPS client fails.

Specifies that WebSphere Application Server challenges the Web client for a user ID and password when the required HTTPS client certificate authentication fails.

The enterprise bean authentication is performed by the Enterprise JavaBean (EJB) authentication module.

The EJB authentication module resides in the CSIv2 and SAS layer.

The authentication module is implemented using the Java Authentication and Authorization Service (JAAS) login module. The Web authenticator and the EJB authenticator pass the authentication data to the login module, which can use the following mechanisms to authenticate the data:

- Kerberos
- LTPA
- RSA token

The authentication module uses the registry that is configured on the system to perform the authentication. Four types of registries are supported:

- Federated repositories
- Local operating system
- Standalone Lightweight Directory Access Protocol (LDAP) registry
- Standalone custom registry

External registry implementation following the registry interface that is specified by IBM can replace either the local operating system or the LDAP registry.

The login module creates a JAAS subject after authentication and stores the credential that is derived from the authentication data in the public credentials list of the subject. The credential is returned to the Web authenticator or to the enterprise beans authenticator.

The Web authenticator and the enterprise beans authenticator store the received credentials in the Object Request Broker (ORB) current for the authorization service to use in performing further access control checks. If the credentials are forwardable, they are sent to other application servers.

You can configure authentication mechanisms in the administrative console by doing the following:

1. **Click Security > Global security.**
2. Under Authentication mechanisms and expiration, select an authentication mechanism to configure.

Lightweight Third Party Authentication

Lightweight Third Party Authentication (LTPA) is intended for distributed, multiple application server and machine environments. LTPA supports forwardable credentials and single sign-on (SSO). LTPA can support security in a distributed environment through cryptography. This support permits LTPA to encrypt, digitally sign, and securely transmit authentication-related data, and later decrypt and verify the signature.

Application servers distributed in multiple nodes and cells can securely communicate using the LTPA protocol. It also provides the single sign-on (SSO) feature wherein a user is required to authenticate only once in a domain name system (DNS) domain and can access resources in other WebSphere Application Server cells without getting prompted. The realm names on each system in the DNS domain are case sensitive and must match identically.

UNIX For local OS, the realm name is the same as the host name.

Windows For local OS, the realm name is the domain name, if a domain is in use or the realm name is the machine name.

For the Lightweight Directory Access Protocol (LDAP), the realm name is the host:port value of the LDAP server.

The LTPA protocol uses cryptographic keys to encrypt and decrypt user data that passes between the servers. These keys must be shared between the different cells for the resources in one cell to access resources in other cells, assuming that all the cells involved use the same LDAP or custom registry.

When using LTPA, a token is created with the user information and an expiration time and is signed by the keys. The LTPA token is time sensitive. All product servers that participate in a protection domain must have their time, date, and time zone synchronized. If not, LTPA tokens appear prematurely expired and cause authentication or validation failures.

This token passes to other servers, in the same cell or in a different cell through cookies, for Web resources when SSO is enabled, or through the authentication protocol layer for enterprise beans.

If the receiving servers share the same keys as the originating server, the token can be decrypted to obtain the user information, which then is validated to make sure that it has not expired and that the user information in the token is valid in its registry. On successful validation, the resources in the receiving servers are accessible after the authorization check.

Each server must have valid credentials. When the credentials expire, the server is required to communicate to the user registry to authenticate. User registry outages can cause server processes to hang, requiring them to be restarted to recover. Extending the time the LTPA token remains cached reduces this risk, but does present a slightly increased security risk to be considered when defining your security policies.

All of the WebSphere Application Server processes in a cell (deployment manager, nodes, application servers) share the same set of keys. If key sharing is required between different cells, export them from one cell and import them to the other. For security purposes, the exported keys are encrypted with a random generated key and a user-defined password is used to protect the keys. This same password is needed when importing the keys into another cell. The password is only used to protect the keys and is not used to generate the keys.

WebSphere Application Server supports the LTPA and Kerberos protocols.

When security is enabled during profile creation time, LTPA is configured by default.

LTPA requires that the configured user registry be a centrally shared repository such as LDAP or a Windows domain-type registry so that users and groups are the same, regardless of the machine.

Lightweight Third Party Authentication key sets and key set groups

Key set groups contain lists of key sets and Lightweight Third Party Authentication (LTPA) key generation schedules. Each key set contains key references to keys in key stores. To generate keys automatically, each key set must be a member of a key set group.

The keys for some key configurations must be generated together. The LTPA key pair is referenced in one key set while the secret or private key is in a separate key set. When the key set group is created, the two key sets are added as members of the key set group. Key set group settings determine whether the keys for both key sets are generated together automatically or manually.

The key set group contains the following attributes:

- Member key sets
- Choice of either manual or automatic key generation in the member key sets
- Schedule for automatically generating keys

Configuring the Lightweight Third Party Authentication mechanism

You must configure Lightweight Third Party Authentication (LTPA) or Kerberos when you set up security for the first time.

1. Open the administrative console.

Type `http://fully_qualified_host_name:port_number/ibm/console` to access the administrative console in a Web browser.

Port 9060 is the default port number for accessing the administrative console. During installation, however, you might have specified a different port number. Use the appropriate port number.

2. Click **Security > Global security > Authentication mechanisms and expiration**.
3. Click **LTPA**.
4. Select the appropriate group from the **Key set group** field that contains your public, private, and shared LTPA keys. These keys are used to encrypt and decrypt data that is sent between servers. You

can access these key set group configurations using the Key set group link. In the Key set group configuration, you can indicate whether to automatically generate new keys and when to generate them.

5. Enter a positive integer value in the **Authentication cache timeout** field. This timeout value refers to how long an LTPA token is valid in minutes. The token contains this expiration time so that any server that receives the token can verify that the token is valid before proceeding further. When the token expires, the user must log in again. An optimal value for this field depends on your configuration. However, the default value is 10 minutes.
6. Enter a positive integer in the **Timeout value for forwarded credentials between servers** field. This value refers to how long the server credentials from another server are valid before they expire. The default value is 120 minutes. The value in the **Timeout value for forwarded credentials between servers** field must be greater than the value in the **Authentication cache timeout** field.
7. Enter a password in the **Password** field, which is used to protect the generated keys that are used to encrypt and decrypt the LTPA keys from the SSO properties file. The password is not used to generate keys, only to protect them. During import, this password should match the password used to export the keys at another LTPA server (for example, another application server Cell, Lotus Domino Server, and so on). During export, remember this password in order to provide it during the import operation. Single sign-on across cells can be provided by sharing keys and passwords. To share the keys and password, log on to one cell, specify a key file, and click **Export keys**. Then, log on to the other cell, specify the key file, and click **Import keys**.
8. Click **Apply** or **OK**. The LTPA configuration is now set. Do not generate the LTPA keys in this step because they are automatically generated later. Proceed with the rest of the steps that are required to enable security, and start with single sign-on (SSO), if it is required.
9. Complete the information in the **Security > Global security** panel and click **OK**. The LTPA keys are generated automatically the first time. Do not generate the keys manually.

Results

The previous steps configured LTPA.

What to do next

After configuring LTPA, you can also complete the following tasks:

1. Generate key files. For more information, see “Generating Lightweight Third Party Authentication keys” on page 254.
2. Export key files. For more information, see “Exporting Lightweight Third Party Authentication keys” on page 255.
3. Import key files. For more information, see “Importing Lightweight Third Party Authentication keys” on page 256.
4. Manage LPTA keys from multiple cells. For more information, see `tsec_ssldmanagetakeys.dita`.
5. If you are enabling security, you can also enable single sign-on (SSO). See:
 - “Configuring single sign-on capability with Tivoli Access Manager or WebSEAL” on page 354
6. If you generated a new set of keys or imported a new set of keys, verify that the keys are saved to the master configuration by clicking **Save** at the top of the panel. Because LTPA authentication uses time-sensitive tokens, verify that the time, date, and time zone are synchronized among all of the product servers that are participating in the protected domain. Changes to the time, date, and time zone are done independently from WebSphere Application Server. If the clock skew is too high between servers, the LTPA token seems prematurely expired and causes authentication or validation failures.

Authentication mechanisms and expiration

Use this page to specify the shared keys and configure the authentication mechanism that is used to exchange information between servers. You can also use this page to specify the amount of time that the authentication information remains valid and specify the single sign-on configuration.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Authentication mechanisms and expiration > LTPA**.

After you configure the properties on this page, complete the following steps:

1. Click **Security > Global security**.
2. Under Available realm definitions, verify that the appropriate registry is configured.
3. Click **Apply**. When security is enabled and any of these properties change, return to the Global security panel and click **Apply** to validate the changes.

Key set group:

Specifies groups of public, private, and shared keys. These key groups enable the application server to manage multiple sets of Lightweight Third Party Authentication (LTPA) keys.

Generate Keys:

Specifies whether to generate a new set of LTPA keys in the configured keystore, and to update the runtime with the new keys. By default, LTPA keys are regenerated on a schedule every 90 days, configurable to the day of the week.

Each new set of LTPA keys is stored in the keystore associated with the key set group. A maximum number of keys (or even one) can be configured. However, it is recommended to have at least two keys; the old keys can be used for validation while the new keys are being distributed.

This step is not necessary during security enablement. A default set of keys is created during the first server startup. If any nodes are down during a key generation event, the nodes should be synchronized with the Deployment Manager before restart.

Authentication cache timeout:

Specifies the time period at which the authenticated credential in the cache expires. Verify that this time period is less than the value for the Timeout value for forwarded credentials between servers field.

If the application server infrastructure security is enabled, the security cache timeout can influence performance. The timeout setting specifies how often to refresh the security-related caches. Security information pertaining to beans, permissions, and credentials is cached. When the cache timeout expires, all cached information not accessed within the timeout period is purged from the cache. Subsequent requests for the information result in a database lookup. Sometimes, acquiring the information requires invoking a Lightweight Directory Access Protocol (LDAP)-bind or native authentication. Both invocations are relatively costly operations for performance. Determine the best trade off for the application, by looking at usage patterns and security needs for the site.

The default security cache timeout value is 10 minutes. If you have a small number of users, it should be set higher than that, or if a large number of users, it should be set lower.

The LTPA timeout value should not be set lower than the security cache timeout. It is also recommended that the LTPA timeout value should be set higher than the orb request timeout value. However, there is no relation between the security cache timeout value and the orb request timeout value.

In a 20-minute performance test, setting the cache timeout so that a timeout does not occur yields a 40% performance improvement.

Data type	Integer
Units	Minutes and seconds

Default 10 minutes
Range: Greater than 30 seconds

Timeout value for forwarded credentials between servers:

Specifies the period of time after which forwarded credentials expire.

Specify a value for this field that is greater than the authentication cache timeout value.

Default 120 minutes

Note: The Timeout value for forwarded credentials between servers should be an integer between the range of 5 through 35971.

Password:

Enter a password which will be used to encrypt and decrypt the LTPA keys from the SSO properties file. During import, this password should match the password used to export the keys at another LTPA server (for example, another application server Cell, Lotus Domino Server, and so on). During export, remember this password in order to provide it during the import operation.

After the keys are generated or imported, they are used to encrypt and decrypt the LTPA token. Whenever the password is changed, a new set of LTPA keys are automatically generated when you click **OK** or **Apply**. The new set of keys is used after the configuration changes are saved.

Data type String

Confirm password:

Specifies the confirmed password that is used to encrypt and decrypt the LTPA keys.

Use this password when importing these keys into other application server administrative domain configurations and when configuring SSO for a Lotus Domino server.

Data type String

Fully qualified key file name:

Specifies the name of the file that is used when importing or exporting keys.

Enter a fully qualified key file name, and click **Import Keys** or **Export Keys**.

Data type String

Internal server ID:

Specifies the server ID that is used for interprocess communication between servers. The server ID is protected with an LTPA token when sent remotely. You can edit the internal server ID to make it identical to server IDs across multiple application server administrative domains (cells). By default this ID is the cell name.

This internal server ID should only be used in a Version 6.1 or higher environment. For mixed-version Cells, you should convert to using a server user ID and server password for interoperability.

To switch back to the server user ID and password for interoperability, complete the following steps:

1. Click **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select a user registry, and click **Configure**.
3. Select the **Server identity that is stored in the repository** option and type a valid registry ID and password.

Data type String

Import Keys:

Specifies whether the server imports new LTPA keys.

To support single sign-on (SSO) in the application server product across multiple application server domains (cells), share the LTPA keys and the password among the domains. You can use the **Import Keys** option to import the LTPA keys from other domains. The LTPA keys are exported from one of the cells to a file. To import a new set of LTPA keys, complete the following steps:

1. Enter the appropriate password in the Password and Confirm password fields.
2. Click **OK** and click **Save**.
3. Enter the directory location where the LTPA keys are located in the Fully qualified key file name field prior to clicking **Import keys**.
4. Do not click **OK** or **Apply**, but save the settings.

Export Keys:

Specifies whether the server exports LTPA keys.

To support single sign-on (SSO) in the WebSphere product across multiple application server domains (cells), share the LTPA keys and the password among the domains. Use the Export Keys option to export the LTPA keys to other domains.

To export the LTPA keys, make sure that the system is running with security enabled and is using LTPA. Enter the file name in the Fully qualified key file name field and click **Export Keys**. The encrypted keys are stored in the specified file.

Generating Lightweight Third Party Authentication keys

WebSphere Application Server generates Lightweight Third Party Authentication (LTPA) keys automatically during the first server startup. You can generate additional keys as you need them in the Authentication mechanisms and expiration panel.

Before you begin

At runtime, the default key sets are CellLTPASecret and CellLTPAKeyPair. The default key group is CellLTPAKeySetGroup. After generation, keys are stored in the default key store CellLTPAKeys.

About this task

Complete the following steps to generate new LTPA keys in the administrative console.

1. Access the administrative console.
Type `http://fully_qualified_host_name:port_number/ibm/console` to access the administrative console in a Web browser.
2. Verify that all the WebSphere Application Server processes are running, including the cell, nodes, and application servers.

Note: If any of the servers are down at the time of key generation and then restarted later, these servers might contain old keys. Copy the new set of keys to these servers to restart them after you generate them.

3. Click **Security > Global security > Authentication mechanisms and expiration**.
4. Click **LTPA**.
5. Click **Generate keys** to generate a new set of LTPA keys in the local keystore and update the runtime with the new keys. By default, LTPA keys are regenerated on a schedule every 90 days, configurable to the day of the week. Each new set of LTPA keys is stored in the keystore that is associated with the key set group. The same password that is already stored in the configuration is used when you generate new keys.

Note: This step is not necessary when you enable security because, by default, a set of keys is created during the first server startup. However, the keystore should have at least two keys: the old keys can be used for validation while the new keys are being distributed. If any nodes are down during a key generation event, the nodes should be synchronized with the Deployment Manager before restarting the server.

6. Restart the server for the changes to become active.

Results

After WebSphere Application Server generates and saves a new set of keys, the generated keys are not used in the configuration until WebSphere Application Server is restarted. Token generation uses the keys that were last imported. To view the latest key version, see “Changing the number of active LTPA keys” on page 259.

What to do next

You must recycle the node agents and application servers to accept the new keys. If any of the node agents are down, run a manual file synchronization utility from the node agent machine to synchronize the security configuration from the deployment manager.

Exporting Lightweight Third Party Authentication keys

To support single sign-on (SSO) in WebSphere Application Server across multiple WebSphere Application Server domains or cells, you must share the Lightweight Third Party Authentication (LTPA) keys and the password among the domains.

Before you begin

Make sure that the time in the domains is similar so that you do not mistakenly interpret the tokens as expired between the cells.

About this task

Complete the following steps in the administrative console to export key files for LTPA so that they can be shared across domains:

1. Type `http://server_name:port_number/ibm/console` in a Web browser to access the administrative console.
2. Click **Security > Global security > Authentication mechanisms and expiration**.
3. Click **LTPA**.
4. In the Password and Confirm password fields, enter the password that is used to encrypt the LTPA keys. Remember the password so that you can use it later when the keys are imported into the other cell.
5. In the Fully qualified key file name field, specify the fully qualified path to the location where you want the exported LTPA keys to reside. You must have write permission to this file.

6. Click **Export keys** to export the keys to the location that you specified in the **Fully qualified key file name** field.
7. Specify the **Internal server ID** that is used for interprocess communication between servers. The server ID is protected with an LTPA token when sent remotely. You can edit the internal server ID to make it identical to server IDs across multiple application server administrative domains (cells). By default this ID is the cell name.
8. Click **OK** and **Save**.

Results

You can share LTPA keys and passwords among domains on WebSphere Application Server.

What to do next

After exporting the keys from one cell, you must import those keys into the other cell. For more information, see “Importing Lightweight Third Party Authentication keys”

Importing Lightweight Third Party Authentication keys

To support single sign-on (SSO) in WebSphere Application Server across multiple WebSphere Application Server domains or cells, you must share the LTPA keys and the password among the domains. You can import LTPA keys from other domains and export keys to other domains.

Before you begin

After you export LTPA keys from one cell, you must import these keys into another cell. To import keys, you must know the password for the exported key file to access the LTPA keys. Verify that key files are exported from one of the cells into a file.

About this task

Complete the following steps in the administrative console to import key files for LTPA.

1. Access the administrative console for the cell that will receive the imported keys by typing `http://server_name:port_number/ibm/console` in a Web browser.
2. Click **Security > Global security > Authentication mechanisms and expiration**.
3. Click **LTPA**.
4. In the **Password** and **Confirm password** fields, enter the password that is used to decrypt the LTPA keys . This password must match the password that was used in the cell from which you are importing the keys.
5. In the **Fully qualified key file name** field, specify the fully qualified path to the location where the signer keys reside. You must have write permission to this file.
6. Click **Import keys** to import the keys to the location that you specified in the **Fully qualified key file name** field.
7. Click **OK** and **Save** to save the changes to the master configuration. It is important to save the new set of keys to match the new password so that no problems are encountered when starting the servers later.

What to do next

After a new set of keys is generated and saved, the generated keys are not used in the configuration until WebSphere Application Server is restarted.

You must recycle the node agents and application servers to accept the new keys. If any of the node agents are down, run a manual file synchronization utility from the node agent machine to synchronize the security configuration from the deployment manager.

Note: After you enter the password in the Password and Confirm password fields and click **Save**, the password is not redisplayed on the administrative console panel.

Disabling automatic generation of Lightweight Third Party Authentication keys

You can disable the automatic generation of new Lightweight Third Party Authentication (LTPA) keys for key sets that are members of a key set group. Automatic generation creates new keys on a schedule that you specify when you configure a key set group, which manages one or more key sets. WebSphere Application Server uses key set groups to automatically generate cryptographic keys or multiple synchronized key sets.

Before you begin

You must know the name of the key set group and the management scope where the key set group is defined.

The default key set group is CellLTPAKeySetGroup.

About this task

LTPA keys are used to encrypt the LTPA token.

Note: You might want to disable the automatic generation of these keys so that you can generate them on a schedule. You should definitely disable automatic key generation if you disable node automatic synchronization. This disabling eventually causes the LTPA keys to fall out of synchronization between the deployment manager and the node agents. Also, you should disable automatic key generation if you import or export LTPA keys to or from another cell. The automatic generation of LTPA keys changes keys over time and causes the cells to fall out of synchronization.

The following steps are needed to complete this task in the administrative console.

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations**.
2. Expand the tree to the inbound or outbound management scope that contains the key set group, and then click the scope link.
3. Under Related Items, click **Key Set Groups**.
4. Click the key set group that you want to disable.
5. Clear the **Automatically generate keys** option.
6. Click **OK** and **Save** to save the changes to the master configuration.
7. Start the server again for the changes to become active.

Results

You have disabled the automatic generation of LTPA keys for the key sets in the key set group.

Note: You can generate keys manually at any time by completing the following steps:

1. Open the key set group collection.
2. Select the check box beside the key set group.
3. Click **Generate keys**.

Managing LTPA keys from multiple WebSphere Application Server cells

You can specify the shared keys and configure the authentication mechanism that is used to exchange information between servers to import and export LTPA keys across multiple WebSphere Application Server cells.

Before you begin

You must be sure that the exported key file for the multiple cells is accessible on the host where WebSphere Application Server is running. Also, you must know the password that was used when the keys were exported.

Note: You should disable automatic key generation if you import or export keys to or from another cell. This disabling causes the imported keys to get lost and the exported keys to no longer interoperate with this cell over time.

At runtime, the default key sets are CellLTPASecret and CellLTPAKeyPair. The default key group is CellLTPAKeySetGroup. After generation, keys are stored in the default key store CellLTPAKeys.

About this task

Complete the following steps to manage LTPA keys using the administrative console.

1. Access the administrative console.
Type `http://fully_qualified_host_name:port_number/ibm/console` to access the administrative console in a Web browser.
2. Verify that all of the WebSphere Application Server processes are running, including cells, nodes, and all of the application servers. If any of the servers are down at the time of key generation and then brought back up later, these servers might contain old keys. Copy the new set of keys to these servers, then bring them back up.
3. Click **Security > Global security > Authentication mechanisms and expiration**.
4. Click **LTPA**.
5. Type the password for the LTPA keys in the **Password** field. Enter a password that is used to encrypt and decrypt the LTPA keys from the single sign-on (SSO) properties file. During import, this password should match the password that is used to export the keys at another LTPA server. During export, remember this password in order to provide it during the import operation.
6. Type the password again in the **Confirm password** field.
7. Select from among the following options:
 - To support SSO in the WebSphere product across multiple application server domains (cells), you can share the LTPA keys and the password among the domains. Before exporting, make sure that security is enabled and using LTPA on the system that is running. For more information, see “Exporting Lightweight Third Party Authentication keys” on page 255.
 - To support SSO in the application server product across multiple application server domains (cells), you can share the LTPA keys and the password among the domains. For more information, see “Importing Lightweight Third Party Authentication keys” on page 256.
 - To import LTPA keys for the current cell if they were previously exported, see “Importing Lightweight Third Party Authentication keys” on page 256.
8. Start the server again for any changes you make to become active.

Results

The shared LTPA keys are now available for WebSphere Application Server to use for secure connections.

What to do next

After the keys are generated or imported, they are used to encrypt and decrypt the LTPA token. To view the latest key version, see “Changing the number of active LTPA keys” on page 259.

Changing the number of active LTPA keys

Key sets manage Lightweight Third Party Authentication (LTPA) keys in a key store that is based on a key alias prefix. A key alias prefix is automatically generated when you generate a new key and store it in a key store. Key stores can contain multiple versions of keys for any given key alias prefix. You can specify a maximum number of active keys in the key set configuration.

Before you begin

You must know the name of the key set group and the management scope where the key set group is defined.

The default key set group is CellLTPAKeySetGroup.

Complete the following steps in the administrative console.

About this task

LTPA keys are used to encrypt the LTPA token. You might want to set a specific number of active keys that WebSphere Application Server returns when the server queries for keys for a particular key set. The following steps are needed to complete this task in the administrative console.

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations**.
2. Expand the tree to the inbound or outbound management scope that contains the key set group, and then click the scope link.
3. Under Related Items, click **Key Sets**.
4. Click the key set that you want to modify.
5. In the **Maximum number of keys referenced** field, type a numerical value for the maximum number of keys that you want to activate.
6. Click **OK** and **Save** to save the changes to the master configuration.
7. Start the server again for the changes to become active. WebSphere Application Server activates only the number of recent keys that you specified.

Results

The **Maximum number of keys referenced** value determines how many active keys are returned when the server queries for keys for the selected key set.

What to do next

You can click **Active key history** in the Key set panel to display the keys that are active for this key set.

Kerberos (KRB5) authentication mechanism support for security

The Kerberos authentication mechanism enables interoperability with other applications (such as .NET, DB2 and others) that support Kerberos authentication. It provides single sign on (SSO) end-to-end interoperable solutions and preserves the original requester identity.

Note: Security support for Kerberos as the authentication mechanism has been added for this release of WebSphere Application Server. Kerberos is a mature, flexible, open, and very secure network authentication protocol. Kerberos includes authentication, mutual authentication, message integrity and confidentiality and delegation features.

The following sections describe Kerberos authentication in more detail:

- “What is Kerberos?” on page 260

- “The benefits of having Kerberos as an authentication mechanism”
- “Kerberos authentication in a single Kerberos realm environment”
- “Kerberos authentication in a cross or trusted Kerberos realm environment” on page 262
- “Setting up Kerberos as the authentication mechanism for WebSphere Application Server” on page 265
- “Setting up Kerberos as the authentication mechanism for the pure Java client” on page 266

What is Kerberos?

Kerberos has withstood the test of time and is now at version 5.0. Kerberos enjoys wide spread platform support (for example, for Windows, Linux, Solaris, AIX, and z/OS) partly because the Kerberos source code is freely downloadable from the Massachusetts Institute of Technology (MIT) where it was originally created.

Kerberos is composed of three parts: a client, a server, and a trusted third party known as the Kerberos Key Distribution Center (KDC). The KDC provides authentication and ticket granting services.

The KDC maintains a database or repository of user accounts for all of the security principals in its realm. Many Kerberos distributions use file-based repositories for the Kerberos principal and policy DB and others use Lightweight Directory Access Protocol (LDAP) as the repository.

Kerberos does not support any notion of groups (that is, iKeys groups or groups of users or principals). The KDC maintains a long-term key for each principal in its accounts database. This long-term key is derived from the password of the principal. Only the KDC and the user that the principal represents should know what the long-term key or password is.

The benefits of having Kerberos as an authentication mechanism

The benefits of having Kerberos as the authentication mechanism for WebSphere Application Server include the following:

- The Kerberos protocol is a standard. This enables interoperability with other applications (such as .NET, DB2 and others) that support Kerberos authentication. It provides single sign on (SSO) end-to-end interoperable solutions and preserves the original requester identity.
- When using Kerberos authentication, the user clear text password never leaves the user machine. The user authenticates and obtains a Kerberos ticket granting ticket (TGT) from a KDC by using a one-way hash value of the user password. The user also obtains a Kerberos service ticket from the KDC by using the TGT. The Kerberos service ticket that represents the client identity is sent to WebSphere Application Server for authentication.
- A Java client can participate in Kerberos SSO using the Kerberos credential cache to authenticate to WebSphere Application server.
- J2EE, Web Service, .NET and Web browser clients that use the HTTP protocol can use the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) token to authenticate to the WebSphere Application Server and participate in SSO by using SPNEGO Web authentication. Support for SPNEGO as the Web authentication service is new to this release of WebSphere Application Server.
Read about “Single sign-on for HTTP requests using SPNEGO Web authentication” on page 302 for more information.
- WebSphere Application Server can support both Kerberos and Lightweight Third-Party Authentication (LTPA) authentication mechanisms at the same time.
- Server-to-server communication using Kerberos authentication is provided.

Kerberos authentication in a single Kerberos realm environment

WebSphere Application Server supports Kerberos authentication in a single Kerberos realm environment as shown in the following figure:

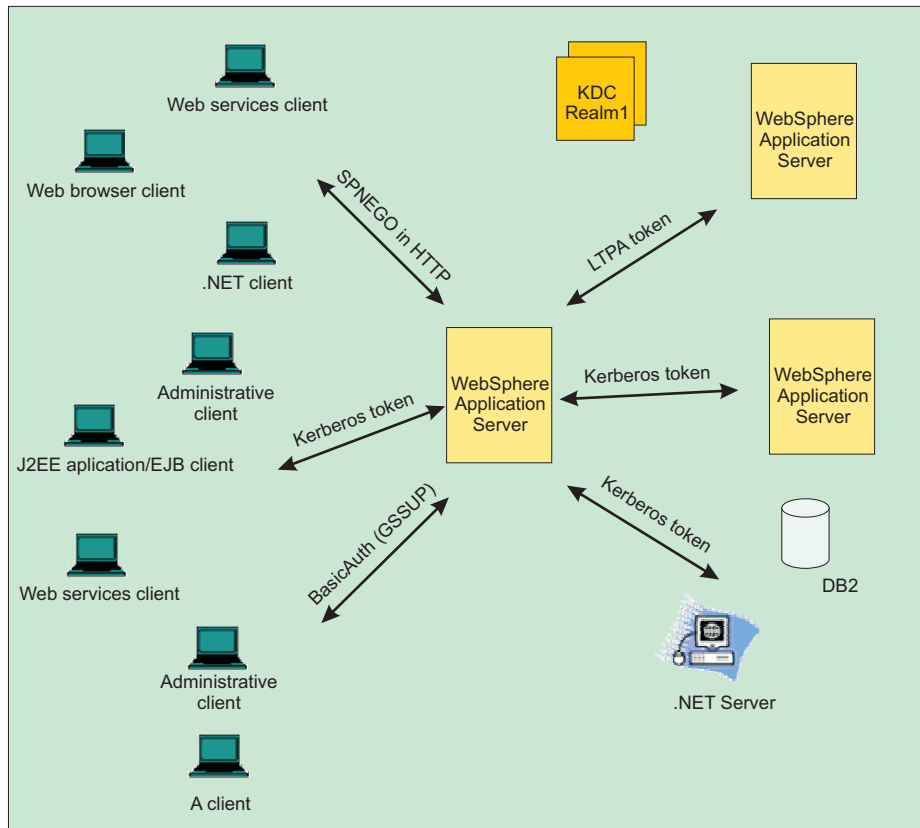


Figure 5. Kerberos authentication in a single Kerberos realm environment

WebSphere Application Server now supports SPNEGO tokens in the HTTP header, Kerberos tokens, LTPA tokens and BasicAuth (GSSUP) for authentication. However, to provide end-to-end Kerberos and end-to-end SPNEGO to Kerberos solutions the following must occur:

- A client must obtain a ticket-granting ticket (TGT) with forwardable, address-less and renewable flags. The renewable flag is optional.
- The **Enabled delegation of Kerberos credentials** option must be selected. Read about “Configuring Kerberos as the authentication mechanism using the administrative console” on page 266 for more information about this option.

When the WebSphere Application Server receives a Kerberos or SPNEGO token for authentication, it uses the Kerberos service principal (SPN) to establish a security context with a requestor. If a security context is established, the WebSphere Kerberos login module extracts a client GSS delegation credential, creates a Kerberos authentication token base on the Kerberos credential, and places them in the client subject with other tokens.

If the server must use a downstream server or back-end resources, it uses the client GSS delegation credential. If a downstream server does not support Kerberos authentication, the server uses the LTPA token instead of the Kerberos token. If a client does not include a GSS delegation credential in the request, the server uses the LTPA token for the downstream server. The Kerberos authentication token and principal are propagated to the downstream server as part of the security attributes propagation feature.

If the WebSphere Application Server and the KDC do not use the same user registry, then a JAAS custom login module might be required to map the Kerberos principal name to the WebSphere user name.

Kerberos authentication in a cross or trusted Kerberos realm environment

WebSphere Application Server also supports Kerberos authentication in a cross or trusted Kerberos realm environment as shown in the following figure:




Figure 6. Kerberos authentication in a cross or trusted Kerberos realm environment

WebSphere Application Server now supports SPNEGO tokens in the HTTP header, Kerberos tokens, LTPA tokens and BasicAuth (GSSUP) for authentication. However, to provide end-to-end Kerberos and end-to-end SPNEGO to Kerberos solutions the following must occur:

- A client must obtain a ticket-granting ticket (TGT) with forwardable, address-less and renewable flags. The renewable flag is optional.
- The **Enabled delegation of Kerberos credentials** option must be selected. Read about “Configuring Kerberos as the authentication mechanism using the administrative console” on page 266 for more information about this option.

When the WebSphere Application Server receives a Kerberos or SPNEGO token for authentication, it uses the Kerberos service principal (SPN) to establish a security context with a requestor. If a security context is established, the WebSphere Kerberos login module always extracts a client GSS delegation credential and Kerberos ticket and places them in the client subject with other tokens.

If the server must use a downstream server or backend resources, it uses the client GSS delegation credential. If a downstream server does not support Kerberos authentication, the server uses the LTPA token instead of the Kerberos token. If a client does not include a GSS delegation credential in the request, the server uses the LTPA token for the downstream server. The Kerberos authentication token and principal are propagated to the downstream server as part of the security attributes propagation feature.

If the WebSphere Application Server and the KDC do not use the same user registry, then a JAAS custom login module might be required to map the Kerberos principal name to the WebSphere user name.

In this release of WebSphere Application Server, the new security multiple domains only support Kerberos at the cell level. All WebSphere application servers must be used by the same Kerberos realm. However, the clients and or backend resources (such as DB2, .NET server, and others) that support Kerberos authentication can have their own Kerberos realm. Only peer-to-peer and transitive trust cross-realm authentication are supported. The following steps must be performed for trusted Kerberos realms:

- The Kerberos trusted realm setup must be done on each of the Kerberos KDCs. See your Kerberos Administrator and User's guide for more information about how to set up a Kerberos trusted realm.
- The Kerberos configuration file might need to list the trusted realm.
- Add Kerberos trusted realms in the administrative console by clicking **Global security > CSiv2 outbound communications > Trusted authentication realms - outbound**.

The following figure shows a Java and administrative client that uses a Kerberos credential cache to authenticate to WebSphere Application Server with a Kerberos token in a trusted Kerberos realm:

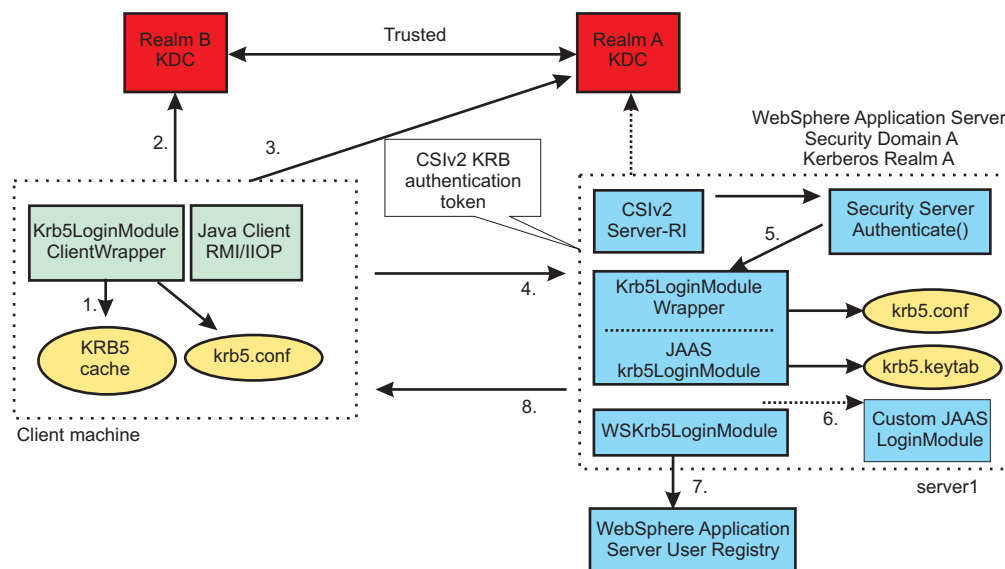


Figure 7. Using a Kerberos credential cache to authenticate to WebSphere Application Server with a Kerberos token in a trusted Kerberos realm

In the figure above, the following events occur:

1. The client uses the Kerberos credential cache if it exists.
2. The client requests a cross realm ticket (TGS_REQ) for Realm A from the Realm B KDC using the Kerberos credential cache.
3. The client uses a cross realm ticket to request Kerberos service ticket for server1 (TGS_REQ) from the Realm A KDC.
4. The Kerberos token returned from the KDC (TGS_REP) is added to the CSiv2 message authentication token and sent to server1 for authentication.
5. The server calls Krb5LoginModuleWrapper to establish security context with the client using the server Kerberos Service Principal Name (SPN) and keys from the **krb5.keytab** file. If the server successfully establishes a security context with the client, it always extracts the client GSS delegation credential and tickets and places them in the client subject.
6. Optionally, a custom JAAS Login Module might be needed if the KDC and WebSphere Application server do not use the same user registry.

7. The user is validated with the WAS user registry.
8. The results (success or failure) are returned to the client.

The following figure shows a Java and administrative client that uses a Kerberos principal name and password to authenticate to WebSphere Application server with a Kerberos token:

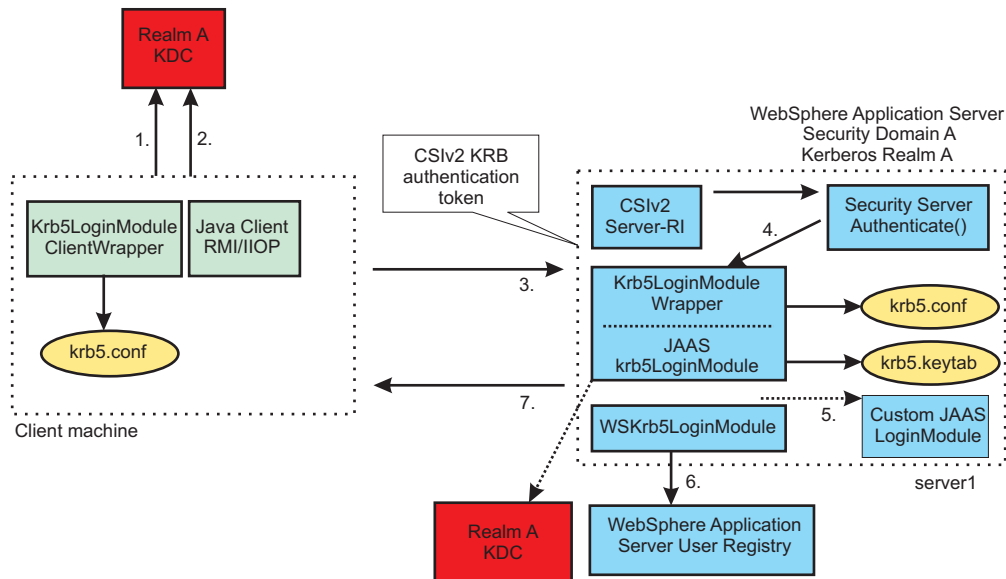


Figure 8. Using a Kerberos principal name and password to authenticate to WebSphere Application Server with a Kerberos token

In the figure above, the following events occur:

1. The client obtains the Kerberos granting ticket (TGT) from the KDC.
2. The client obtains a Kerberos service ticket for server1 (TGS_REQ) using the TGT.
3. The Kerberos token returned from the KDC (TGS_REP) is added to the CSiv2 message authentication token and sent to server1 for authentication.
4. The server calls `Krb5LoginModuleWrapper` to establish security context with the client using the server Kerberos Service Principal Name (SPN) and keys from the `krb5.keytab` file. If the server successfully establishes a security context with the client, it always extracts the client GSS delegation credential and tickets and places them in the client subject.
5. Optionally, a custom JAAS Login Module might be needed if the KDC and WebSphere Application server do not use the same user registry.
6. The user is validated with the WAS user registry.
7. The results are returned to the client.

The following figure shows server-to-server communications:




Figure 9. Server to server communications

When a WebSphere application server starts up, it uses the server ID and password to login to the KDC and then obtains the TGT. It then uses the TGT to request a service ticket to communicate with another server. If a WebSphere application server uses the internal server ID instead of the server ID and password, server-to-server communication is done using an LTPA token. In the figure above, the following events occur:

1. WebSphere Application Server 1 invokes a method, `foo()`, on an Enterprise JavaBeans (EJB) running in WebSphere Application Server 2.
2. Server1 obtains a Kerberos service ticket for Server2 (TGS_REQ) using the Server1 TGT.
3. Same as step 2.
4. The Kerberos token returned from a KDC (TGS_REP) is added to the CSlv2 message authentication token and sent to Server2 for authentication.
5. Server2 calls the `acceptSecContext()` method to establish security context with server1 using the server2 Kerberos Service Principal Name (SPN) and keys from the **krb5.keytab** file. If server2 successfully establishes a security context with server1, it always extracts the server1 GSS delegation credential and tickets and places them in the subject.
6. The server id is validated with the WebSphere user registry.

Setting up Kerberos as the authentication mechanism for WebSphere Application Server

You must perform the following steps in order to set up Kerberos as the authentication mechanism for WebSphere Application Server.

Note: Kerberos authentication mechanism on the server side must be done by the system administrator and on the Java client side by end users. The Kerberos keytab file must to be protected.

1. Ensure that the KDC is configured.
See your Kerberos Administrator and User's guide for more information.
2. The IBM implementation of the Java Generic Security Service (JGSS) and KRB5 require a Kerberos configuration file (**krb5.conf** or **krb5.ini**) on each node or Java virtual machine (JVM). In this release of WebSphere Application Server, this configuration file should be placed in the `config/cells/<cell_name>` directory so that all application servers can access this file. If you do not have a Kerberos configuration file, use a **wsadmin** command to create one.

Read about “Creating a Kerberos configuration file” on page 270 for more information.

3. Kerberos prefers servers and services to have a host-based service ID. The format of this ID is `<service name>/<fully qualified hostname>`. The default service name is WAS. For Kerberos authentication, the service name can be any strings that are allowed by the KDC. However, for SPNEGO Web authentication, the service name must be HTTP. An example of a WebSphere Application Server server ID is `WAS/myhost.austin.ibm.com`.

Each host must have a server ID unique to the hostname. All processes on the same node share the same host-based service ID.

A Kerberos administrator creates a Kerberos service principal name (SPN) for each node in the WebSphere cell. For example, for a cell with 3 nodes (such as `server1.austin.ibm.com`, `server2.austin.ibm.com` and `server3.austin.ibm.com`), the Kerberos administrator must create the following Kerberos service principals: `WAS/server1.austin.ibm.com`, `WAS/server2.austin.ibm.com` and `WAS/server3.austin.ibm.com`.

The Kerberos keytab file (**krb5.keytab**) contains all of the SPNs for the node and must be protected. This file can be placed in the `config/cells/<cell_name>` directory

Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information.

4. Set up Kerberos as the authentication mechanism for WebSphere Application server on the server side by using the administrative console.

Read about “Configuring Kerberos as the authentication mechanism using the administrative console” for more information.

For information about how to set up Kerberos as the authentication mechanism for WebSphere Application server on the server side by using **wsadmin** commands, read about “Kerberos authentication commands” on page 1262.

5. A custom JAAS Login Module might be needed if the KDC and WebSphere Application server do not use the same user registry. Read about “Mapping of a client Kerberos principal name to the WebSphere user registry ID” on page 276 for more information.

Setting up Kerberos as the authentication mechanism for the pure Java client

End users can set up Kerberos authentication mechanism for the pure Java client.

Read about “Configuring a Java client for Kerberos authentication” on page 280 for more information.

Configuring Kerberos as the authentication mechanism using the administrative console

You can use the administrative console to configure Kerberos as the authentication mechanism for the application server. When you have entered and applied the required information to the configuration, the Kerberos service principal name is formed as `<service name>/<fully qualified hostname>@KerberosRealm`, and is used to verify incoming Kerberos token requests.

Before you begin

Read about “Kerberos (KRB5) authentication mechanism support for security” on page 259 for a better understanding of how to set up Kerberos as the authentication mechanism in this version of WebSphere Application Server. You must have completed the steps indicated in this information before you attempt to configure Kerberos as the authentication mechanism using the administrative console.

The following items are required before you attempt to configure Kerberos as the authentication mechanism using the administrative console:

- If you do not already have a Kerberos configuration file (**krb5.ini** or **krb5.conf**), use the **createkrbConfigFile** command task to create the Kerberos configuration file. Read about “Creating a Kerberos configuration file” on page 270 for more information.

- You must have a Kerberos keytab file (**krb5.keytab**) that contains a Kerberos service principal name (SPN), <service name>/<fully qualified hostname>@KerberosRealm, for each machine that run WebSphere application servers. The service name can be anything; the default value is WAS.

For example, if you have two application server machines, host1.austin.ibm.com and host2.austin.ibm.com, the Kerberos keytab file must contain the <service name>/host1.austin.ibm.com and service name>/host2.austin.ibm.com SPNs and their Kerberos keys.

Kerberos will only load and use one keytab file per session. For example, if Kerberos is configured, and you want to use a new keytab file with the same name and location as the previous keytab file, you must first restart the server to use the new keytab file.

If you are configuring Kerberos for the first time, and you accidentally use a bad keytab file, you must unconfigure Kerberos and restart the server before you can configure Kerberos again using a new keytab file.

You must first enable global and application security.

If Kerberos is configured in global security, but you want to configure Simple and Protected GSS-API Negotiation (SPNEGO) on a domain using a different Kerberos realm, you must first use the **klist -m** command to merge existing keytab files into one keytab file. Use that merged keytab file to configure Kerberos and SPNEGO on global and domain security

1. In the administrative console, click **Security > Global security**.
2. From Authentication, click **Kerberos configuration**.
3. Enter your Kerberos service name. By convention, a Kerberos service principal is divided into three parts: the primary, the instance, and the Kerberos realm name. The format of the Kerberos service principal name is <serviceName>/<fully qualified hostName>@KERBEROS_REALM. The service name is the first part of the Kerberos service principal name. For example, in WAS/test.austin.ibm.com@AUSTIN.IBM.COM, the service name is WAS. In this example, the keytab file must have the Kerberos service principal name,WAS/test.austin.ibm.com@AUSTIN.IBM.COM, and its keys.
4. Enter the Kerberos configuration file name with its full path, or click **Browse** to locate it. The Kerberos client configuration file, **krb5.conf** or **krb5.ini**, contains Kerberos configuration information, including the locations of the Key Distribution Centers (KDCs) for the realm of interest. The **krb5.conf** file is the default file name for all platforms except the Windows operating system, which uses the **krb5.ini** file.
5. Optional: Enter the Kerberos keytab file name with its full path, or click **Browse** to locate it. The Kerberos keytab file contains one or more Kerberos service principal names and keys. The default keytab file is **krb5.keytab**. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users. Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information. If you do not specify this parameter, the default keytab in the Kerberos configuration file is used.
6. Enter the name of your Kerberos realm in the **Kerberos realm name** field. In most cases, your realm is your domain name in uppercase letters. If you do not specify this parameter, the default Kerberos realm name in the Kerberos configuration file is used.

For example, a machine with the domain name of test.austin.ibm.com would usually have a Kerberos realm name of AUSTIN.IBM.COM.

Note: The Kerberos realm name for the Microsoft KDC is an uppercase of the Microsoft Domain Controller name.

7. Optional: **Trim Kerberos realm from principal name** is selected by default. You can deselect this option if you want the suffix of the Kerberos principal name to be retained. This option specifies whether the Kerberos login module removes the suffix of the principal user name, starting from the @ that precedes the Kerberos realm name. If this attribute is set to true, the suffix of the principal user name is removed. If this attribute is set to false, the suffix of the principal name is retained. The default value used is true.
8. Optional: **Enable delegation of Kerberos credentials** is selected by default. This option specifies whether the Kerberos delegated credentials are to be stored in the subject by the Kerberos

authentication. This option also enables an application to retrieve the stored credentials and to propagate them to other applications downstream for additional Kerberos authentication with the credential from the Kerberos client.

Note: If this parameter is true, and the runtime cannot extract a client GSS delegation credential, then a warning message is displayed.

9. Click **OK**.

Note: When you select **Apply** or **OK** the Kerberos authentication is automatically tested. If the Kerberos configuration is not complete, a message is displayed that indicates authentication failure.

Results

You have now configured and saved Kerberos as the authentication mechanism for WebSphere Application Server.

What to do next

To enable SPNEGO, click **SPNEGO web authentication enablement** from Related Configuration.

SPNEGO Web authentication and Kerberos authentication use the same Kerberos client configuration and keytab files.

Kerberos authentication

Use this page to configure and to verify Kerberos as the authentication mechanism for the application server.

When you have entered and applied the required information to the configuration, the server principal name is created from the service name, realm name, and host name, and is used to automatically verify authentication to the Kerberos service.

When configured, Kerberos is the primary authentication mechanism. Configure Enterprise JavaBeans (EJB) authentication to resources by accessing the resource references links on the application details panel.

To view this administrative console page, click **Security > Global security**. Under Authentication, click **Kerberos configuration**.

Note: When configuring Kerberos, if the configuration fails with an exception such as in the following example:

```
org.ietf.jgss.GSSEException, major code: 11, minor code: 0 major string: General failure,
unspecified at GSSAPI level minor string: Cannot get credential for
principal service WAS/test@AUSTIN.IBM.COM
```

the principal service must be in the format: <service name>/<fully qualified hostname>@KerberosRealm. In the exception example, the fully qualified hostname is not specified, which is why the failure occurs. For this failure, the hostname of the system is usually obtained from the /etc/hosts file instead of from the Domain Name Server (DNS). On UNIX or Linux systems, if the "hosts": line in the /etc/nsswitch.conf file is configured to first look in the hosts file before it looks in DNS, the Kerberos configuration fails if the hosts file contains an entry for the system that is not the fully qualified hostname.

Kerberos realm name:

The name of your Kerberos realm. In most cases, your realm is your domain name in uppercase letters. For example, a machine with the domain name of `test.austin.ibm.com` typically has a Kerberos realm name of `AUSTIN.IBM.COM`. If this field is left empty, then the realm name specified in the Kerberos configuration file is used.

Data type: String

Kerberos service name:

By convention, a Kerberos service principal is divided into three parts: the primary, the instance, and the Kerberos realm name. The format of the Kerberos service principal name is `service/<fully qualified hostname>@KERBEROS_REALM..`. The service name is the first part of the Kerberos service principal name. For example, in `WAS/test.austin.ibm.com@AUSTIN.IBM.COM`, the service name is `WAS`.

Default: String

Kerberos configuration file with full path:

The Kerberos configuration file, `krb5.conf` or `krb5.ini`, contains client configuration information, including the locations of the Key Distribution Centers (KDCs) for the realm of interest. The `krb5.conf` file is used for all platforms except the Windows operating system, which uses the `krb5.ini` file.

Data type: String

Kerberos keytab file name with full path:

Specifies the Kerberos keytab file name with its full path. You can click **Browse** to locate it. If this field is left empty, then the keytab file name specified in the Kerberos configuration file is used.

Data type: String

Trim Kerberos realm from principal name:

Specifies whether Kerberos removes the suffix of the principal user name, starting from the `@` that precedes the Kerberos realm name. If this attribute is set to `true`, the suffix of the principal user name is removed. If this attribute is set to `false`, the suffix of the principal name is retained. The default value used is `true`.

Default:: Enabled

Enable delegation of Kerberos credentials:

Specifies whether the Kerberos delegated credentials are to be stored in the subject by the Kerberos authentication.

This option also enables an application to retrieve the stored credentials and to propagate them to other applications downstream for additional Kerberos authentication with the credential from the Kerberos client.

Note: If this parameter is `true`, and the runtime cannot extract a client GSS delegation credential, then a warning message is logged.

Default:: Enabled

Creating a Kerberos configuration file

The Kerberos configuration file contains client configuration information, including the locations of Key Distribution Centers (KDCs) for the realms of interest, defaults for the current Kerberos realm and mappings of host names onto Kerberos realms.

Before you begin

About this task

Kerberos configuration settings, the Kerberos key distribution center (KDC) name, and realm settings for both Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) Web authentication and Kerberos authentication are provided in the Kerberos configuration file or through the `java.security.krb5.kdc` and `java.security.krb5.realm` JVM system properties.

If you don't have a Kerberos configuration file (**krb5.ini** or **krb5.conf**), you must first create a Kerberos configuration file and place it in every instance in a cell using a **wsadmin** command. The following are the default Kerberos configuration files and their locations:

- On a Windows platform, the default location is `c:\winnt\krb5.ini`.

Note: if the **krb5.ini** file is not located in the `c:\winnt` directory it might be located in `c:\windows`.

- On a Linux platform, the default location is `/etc/krb5.conf`.
- On other Unix platforms, the default location is `/etc/krb5/krb5.conf`.
- On a z/OS platform, the default location is `/etc/krb5/krb5.conf`.
- On i5/OS platform, the default location is `/etc/krb5/krb5.conf`.

Use the **wsadmin** utility to create a Kerberos configuration file for WebSphere® Application Server:

1. Start the command-line utility by running the **wsadmin** command from the `app_server_root/bin` directory.
2. At the **wsadmin** prompt, enter the following command:

```
$AdminTask help createKrbConfigFile
```

You can use the following parameters with the `createKrbConfigFile` command:

Option	Description
<krbPath>	This parameter is required. It provides the fully qualified file system location of the Kerberos configuration (krb5.ini or krb5.conf) file.
<realm>	This parameter is required. It provides the Kerberos realm name. The value of this attribute is used by SPNEGO to form the Kerberos service principal name for each of the hosts specified with the property <code>com.ibm.ws.security.spnego.SPN<id>.hostName</code> .
<kdcHost>	This parameter is required. It provides the host name of the Kerberos Key Distribution Center (KDC).
<kdcPort>	This parameter is optional. It provides the port number of the Kerberos Key Distribution Center. If this port is omitted, it defaults to 88.
<dns>	This parameter is required. It is a list of default domain name services (DNS), separated by a pipe character, that is used to produce a fully qualified host name. The first one in the list is the default domain name service.
<keytabPath>	This parameter is required. It provides the file system location of the Kerberos keytab path and file name.

Option	Description
<encryption>	This parameter is optional. It identifies the list of supported encryption types, separated by a pipe character. The default value is des-cbc-md5.

The following encryption types are supported:

- des-cbc-md5
- des-cbc-crc
- des3-cbc-sha1
- rc4-hmac
- arcfour-hmac
- arcfour-hmac-md5
- aes128-cts-hmac-sha1-96
- aes256-cts-hmac-sha1-96

Note: Not all of the KDC solutions available today support all of the encryption types listed above. Before you choose an encryption type, ensure that your KDC supports the encryption type that you want to use by consulting your Kerberos Administrator's and User's Guide.

Ensure you have a common encryption type for the Kerberos configuration file, the Kerberos keytab file, the Kerberos service principal name and the Kerberos client. For example, if the Kerberos client uses the RC4-HMAC encryption type, the target server must also support the RC4-HMAC encryption type and the Kerberos configuration file must list RC4-HMAC first in **default_tgt_encetypes** and **default_tkt_encetypes**.

The following is an example of the createKrbConfigFile command:

```
$AdminTask createKrbConfigFile {-krbPath c:/winnt/krb5.ini
                                -realm WSSEC.AUSTIN.IBM.COM
                                -kdcHost host1.austin.ibm.com|host2.austin.ibm.com
                                -dns austin.ibm.com|raleigh.ibm.com
                                -keytabPath c:/winnt/krb5.keytab}
```

This example creates the c:/winnt/krb5.ini file as follows:

```
[libdefaults]
default_realm = WSSEC.AUSTIN.IBM.COM
default_keytab_name = FILE:c:/winnt/krb5.keytab
default_tkt_encetypes = rc4-hmac des-cbc-md5
default_tgs_encetypes = rc4-hmac des-cbc-md5
forwardable = true
renewable = true
noaddresses = true
clockskew = 300
[realms]
WSSEC.AUSTIN.IBM.COM = {
    kdc = host1.austin.ibm.com:88
    kdc = host2.austin.ibm.com:88
    default_domain = austin.ibm.com
}
[domain_realm]
.austin.ibm.com = WSSEC.AUSTIN.IBM.COM
.raleigh.ibm.com = WSSEC.AUSTIN.IBM.COM
```

The createKrbConfigFile command creates a simple Kerberos configuration file. You might have to edit this file as needed, especially when you have a cross or trusted realm environment.

The [domain_realm] section of the Kerberos configuration file is needed for a cross realm environment.

- [domain_realm] - Provides a translation from a domain name or host name to a realm name. The tag name can be a host name or a domain name. Domain names are indicated by a prefix of a period ('.'). The value of the relation is the realm name for that particular host or domain. Host names and domain names should be in lowercase.

If no translation entry applies, the host's realm is considered to be the host name's domain portion converted to uppercase. For example, the following [domain_realm] section maps tech.ibm.com into the TEST.AUSTIN.IBM.COM realm:

```
[domain_realm]
.austin.ibm.com = WSSEC.AUSTIN.IBM.COM
.raleigh.ibm.com = WSSEC.AUSTIN.IBM.COM
```

All other hosts in the austin.ibm.com and .raleigh.ibm.com domains map default to WSSEC.AUSTIN.IBM.COM.

The following example contains more than one Kerberos realm name:

```
[domain_realm]
.ibm.com =AUSTIN.IBM.COM
ibm.com =AUSTIN.IBM.COM
tech.ibm.com =TEST.AUSTIN.IBM.COM
.fubar.org =FUBAR.ORG
```

All other hosts in the ibm.com® domain map by default to the AUSTIN.IBM.COM realm and all hosts in the fubar.org domain map by default to the FUBAR.ORG realm.

Note the entries for the hosts, ibm.com and fubar.org. Without these entries, these hosts would map into the realms COM and ORG, respectively.

For *peer trust cross-realm authentication*, see your Kerberos Administrator's and User's Guide for information about how to set up the trust cross-realm authentication on the KDC.

You must add information about the foreign realm to the realms and domain_realm sections of the Kerberos configuration file. For example:

```
[realms]
AUSTIN.IBM.COM = {
    kdc = kdc.austin.ibm.com:88
    default_domain = austin.ibm.com
}
FUBAR.ORG = {
    kdc = kdc.fubar.org:88
    default_domain = fubar.org
}

[domain_realm]
austin.ibm.com = AUSTIN.IBM.COM
.austin.ibm.com = AUSTIN.IBM.COM
fubar.org = FUBAR.ORG
.fubar.org = FUBAR.ORG
```

In a *transitive trust*, two realms trust each other if they trust the intermediate realms involved in granting a ticket. If each realm involved in granting the service ticket is present in the trust path, then the ticket is trusted. See your Kerberos Administrator's and User's Guide for information about how to configure transitive trust on the KDC.

To set up transitive trust, separate realms with cross-realm authentication must be defined between a and b and between b and c, but not between a and c. With transitive trust, REALMA and REALMC can communicate with each other, but only by going through REALMB.

```
REALMA <-> REALMB <-> REALMC
```

A [capaths] stanza must be added to each **krb5.conf** file.

The **krb5.conf** files on all machines must list all three realms in the [realms] stanza. REALMA must list itself and REALMB and REALMC; REALMB must list itself and REALMA and REALMC; REALMC must list itself and REALMA and REALMB. In the [domain_realm] stanza of the **krb5.conf** files list all three host names and realm names to be able to execute from REALMA to REALMC and from REALMC to REALMA.

```
[capaths]
REALMA.AUSTIN.IBM.COM = {
    REALMB.AUSTIN.IBM.COM = .
    REALMC.AUSTIN.IBM.COM = REALMB.AUSTIN.IBM.COM
}
REALMB.AUSTIN.IBM.COM = {
```

```

REALMC.AUSTIN.IBM.COM = .
REALMA.AUSTIN.IBM.COM = .
}
REALMC.AUSTIN.IBM.COM = {
REALMB.AUSTIN.IBM.COM = .
REALMA.AUSTIN.IBM.COM = REALMB.AUSTIN.IBM.COM
}

```

The **krb5.conf** file permission must be set to 644, which means that you can read and write the file. However, members of the group that the file belongs to and all other users can only read the file.

Since the Kerberos configuration and keytab file are set by the JVM system properties, `java.security.krb5.conf` and `KRB5_KTNAME` respectively, if SPNEGO Web authentication and Kerberos authentication are both enabled you must use the same Kerberos configuration and keytab files for both.

Results

You have now created a Kerberos configuration file.

Creating a Kerberos service principal and keytab file

This task is necessary to process SPNEGO Web or Kerberos authentication requests to WebSphere Application Server. You can create a Kerberos service principal and keytab name using Microsoft Windows, iSeries®, Linux, Solaris, Massachusetts Institute of Technology (MIT) and z/OS KDCs.

About this task

Creating a Kerberos service principal and keytab using Microsoft Windows KDC:

This task is performed on the active directory domain controller machine. Complete the following steps to ensure that the Microsoft Windows 2000 or Windows 2003 Server that is running the active directory domain controller is configured properly to the associated key distribution center (KDC).

1. Create a user account in the Microsoft Active Directory for the WebSphere Application Server.

Click **Start->Programs->Administrative Tools->Active Directory Users and Computers**.

Use the name for the WebSphere Application Server. For example, if the Application Server you are running on the WebSphere Application Server machine is called `myappserver.austin.ibm.com`, create a new user in an Active Directory called `myappserver`.

Make sure that you do not have the computer name `myappserver` under Computers and Domain Controllers. If you already have a computer named `myappserver`, then you must create a different user account name.

- Click **Start -> Programs -> Administrative Tools -> Active Directory Users and Computers->Computers**.
- Click **Programs -> Administrative Tools -> Active Directory Users and Computers->Domain Controllers**.

2. Use the **setspn** command to map the Kerberos service principal name, `<service name>/<fully qualified host name>`, to a Microsoft user account.

The service name for SPNEGO Web authentication must be HTTP, but it can be any strings that are allowed by the KDC for Kerberos authentication.

An example of the **setspn** command usage for SPNEGO Web authentication is as follows:

```

C:\Program Files\Support Tools>
setspn -A HTTP/myappserver.austin.ibm.com myappserver

```

Note: The host name must be a fully-qualified host name.

Note: Make sure that you do not have the same service principle names (SPNs) mapping to more than one Microsoft user account. If you map the same SPN to more than one user account, the web browser client can send an NT LAN manager (NTLM) token instead of a SPNEGO token to WebSphere Application Server.

3. Create the Kerberos keytab file and make it available to WebSphere Application Server. Use the **ktpass** tool from the Windows Server toolkit to create the Kerberos keytab file (krb5.keytab) for the service principal name (SPN).

Note: A Kerberos keytab file contains a list of keys that are analogous to user passwords. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk.

Make sure to use the **ktpass** tool that matches the Windows server level you are using. That is, use the Windows 2000 version for a Windows 2000 Server, or a Windows 2003 version for a Windows 2003 server. The Windows 2003 server version of the **ktpass** tool supports the encryption type, RC4-HMAC, and Single data encryption standard (DES). The Windows 2000 server version of the **ktpass** tool is similar, but different options are necessary for the RC4-HMAC encryption type and single DES. For more information about the **ktpass** tool, see Windows 2003 Technical Reference (Kerberos keytab file and ktpass command).

- For a single DES encryption type:

From a command prompt, run the **ktpass** command as in the following example:

```
ktpass -out c:\temp\myappserver.keytab
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM
-mapUser myappserv
-mapOp set -pass was1edu
-crypto DES-CBC-MD5
+DesOnly
```

Table 4. Using ktpass for a single DES encryption type

Option	Explanation
-out c:\temp\myappserver.keytab	The key is written to this output file.
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM	The concatenation of the user logon name, and the realm must be in uppercase.
-mapUser	The key is mapped to the user, myappserver.
-mapOp	This option sets the mapping.
-pass was1edu	This option is the password for the user ID.
-crypto DES-CBC-MD5	This option uses the single DES encryption type.
+DesOnly	This option generates only DES encryptions.

- For the RC4-HMAC encryption type:

Note: RC4-HMAC encryption is only supported when using a Windows 2003 Server as KDC.

RC4-HMAC encryption is not supported with a Windows 2000 Server as KDC.

From a command prompt, run the **ktpass** command as in the following example:

```
ktpass -out c:\temp\myappserver.keytab
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM
-mapUser myappserver
-mapOp set
-pass was1edu
-crypto RC4-HMAC
```

Table 5. Using ktpass for the RC4-HMAC encryption type

Option	Explanation
-out c:\temp\myappserver.keytab	The key is written to this output file.

Table 5. Using ktpass for the RC4-HMAC encryption type (continued)

Option	Explanation
-princ HTTP/myappserver.austin.ibm.com @WSSEC.AUSTIN.IBM.COM	The concatenation of the user logon name, and the realm must be in uppercase.
-mapUser	The key is mapped to the user, myappserver.
-mapOp	This option sets the mapping.
-pass was1edu	This option is the password for the user ID.
-crypto RC4-HMAC	This option chooses the RC4-HMAC encryption type.

The Kerberos keytab file is created for use with SPNEGO.

Creating a Kerberos service principal and keytab file using iSeries, Linux, Solaris and MIT KDCs:

See your Kerberos implementation documents for the **kadmin**, **kadmin.local addprinc** and **ktadd** commands for more detailed information.

This task is performed on a Linux, Solaris or MIT KDC machine.

1. Create a Kerberos service principal for Kerberos authentication. For example:

```
WAS/testmach.austin.ibm.com
kadmin.local: addprinc WAS/testmach.austin.ibm.com
```

2. Add the newly-created Kerberos service principal, WAS/testmach.austin.ibm.com to a default **krb5.keytab** file. For example:

```
kadmin.local: ktadd WAS/testmach.austin.ibm.com
```

Creating a Kerberos service principal and keytab file using z/OS KDC:

Before Simple and Protected GSS-API Negotiation (SPNEGO) Web authentication and Kerberos authentication can be used, the WebSphere Application Server administrator must first create a Kerberos keytab file on the host that is running WebSphere Application Server.

To create an SPN, do the following:

1. The Kerberos ID (KRBNAME) must be of the form *<service>/<fully qualified system name>*.
2. The following example creates the Kerberos SPN for SPNEGO Web, HTTP/host1.pok.ibm.com:
ALTUSER ASCR1 KERB(KRBNAME(HTTP/host1.pok.ibm.com))
3. Generate the Kerberos key for this user. To generate this key, a password must be associated with this ID. You should not allow this ID to log onto the system. Enter the following two lines whenever a new Kerberos key is required.

Note: The WebSphere or KDC administrator must know this password to create an entry in the keytab file.

```
ALTUSER ASCR1 PASSWORD(was1krb) NOEXPIRED
ALTUSER ASCR1 NOPASSWORD
```

4. Verify that this user has a valid Kerberos segment and a key. For example:

```
LISTUSER ASCR1 KERB NORACF
USER=ASCR1
KERB INFORMATION
-----
KRBNAME= HTTP/host1.pok.ibm.com
KEY VERSION= 001
KEY ENCRYPTION TYPE= DES NODES3 NODESD
```

To create a Kerberos keytab (**krb5.keytab**) file, use the Java Kerberos **ktab** command, *<\$WAS_HOME>/java/bin/ktab*, by doing the following:

- From a command line, type the **ktab -help** command to obtain the proper usage for this command. For example:

```
(host1)CTC03:/PYRSA1/usr/lpp/zWebSphere/V7R1/java/J5.0/bin(189):>ktab -help
Usage: java com.ibm.security.krb5.internal.tools.Ktab [options]
Available options:
  -l      list the keytab name and entries
  -a <principal_name> [password] add an entry to the keytab
  -d <principal_name>      delete an entry from the keytab
  -k <keytab_name>      specify keytab name and path with FILE: prefix
```

- Also from a command line, use the **ktab** command to add the SPN to a default keytab file. For example:

```
(host1)CTC03:/PYRSA1/usr/lpp/zWebSphere/V7R1/java/J5.0/bin(201):>ktab -a
HTTP/host1.pok.ibm.com@LSREALM.POK.IBM.COM ot56prod
Done!
Service key for principal HTTP/host1.pok.ibm.com@LSREALM.POK.IBM.COM saved
```

- Verify that the correct SPN is in the default keytab file. For example:

```
(host1)CTC03:/PYRSA1/usr/lpp/zWebSphere/V7R1/java/J5.0/bin(202):>ktab
1 entries in keytab, name: /etc/skrb/krb5.keytab
  KVNO      Principal
  ----      -
  1          HTTP/host1.pok.ibm.com@LSREALM.POK.IBM.COM
```

Make the keytab file available to WebSphere Application Server. Copy the **krb5.keytab** file from the KDC to the WebSphere Application Server machine at the location specified in the Kerberos configuration file (**krb5.ini** or **krb5.conf**). For example:

```
ftp> bin
ftp> put c:\temp\KRB5_NT_SEV_HST\krb5.keytab
```

Note: A Kerberos keytab configuration file contains a list of keys that are analogous to user passwords. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users.

Use the **validateKrbConfig** command to validate the **krb5.conf** and **krb5.keytab** files. For example:

```
wsadmin>$AdminTask help validateKrbConfig
```

Note: The Kerberos keytab file is shared by Kerberos and SPNEGO Web authentication. It is loaded once and cannot be refreshed.

Results

You have created a Kerberos service principal and keytab file on the KDC that WebSphere Application Server uses to process SPNEGO and or Kerberos authentication requests.

Mapping of a client Kerberos principal name to the WebSphere user registry ID

You can map the Kerberos client principal name to the WebSphere user registry ID for both Simple and Protected GSS-API Negotiation (SPNEGO) Web authentication and Kerberos authentication.

Before you begin

About this task

Use the Java Authentication and Authorization Service (JAAS) custom login module to perform any custom mapping of a client Kerberos principal name to the WebSphere user registry identity. The JAAS custom login module is a plug-in mechanism that is defined for authenticating incoming requests in WebSphere Application Server. If the active authentication mechanism is LTPA, the JAAS custom login module is inserted immediately before the `ltpaLoginModule`.

The JAAS custom login module retrieves a client Kerberos principal name in `javax.security.auth.Subject` using the `subject.getPrincipals(KerberosPrincipal.class)` method, maps the client Kerberos principal name to the WebSphere user registry identity, and inserts the mapping identity in the hash table property, `com.ibm.wsspi.security.cred.userId`. The `wsMapDefaultInboundLoginModule` then uses the mapped identity to create a `WSCredential`.

Note: For SPNEGO Web authentication, the custom login module can also supply the full set of security properties in the `javax.security.auth.Subject` in the `com.ibm.wsspi.security.tai.TAIResult` to fully assert the mapped identity. When the identity is fully asserted, the `wsMapDefaultInboundLoginModule` maps those security properties to a `WSCredential`.

1. Refer to the following example of a custom login module:

```
import java.util.Map;
import java.lang.reflect.Array;
import javax.security.auth.Subject;
import javax.security.auth.callback.*;
import javax.security.auth.login.LoginException;
import javax.security.auth.spi.LoginModule;
import javax.security.auth.kerberos.*;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.AttributeNameConstants;

/**
 *
 * @author IBM Corporation
 * @version 1.0
 * @since 1.0
 */

public class sampleSpnegoMappingLoginModule implements LoginModule {
    /**
     *
     * Constant that represents the name of this mapping module. Whenever this sample
     * code is used to create a class with a different name, this value should be changed.
     */
    private final static String MAPPING_MODULE_NAME = "com.ibm.websphere.security.sampleSpnegoMappingLoginModule";

    private String mapUid = null;
    /**
     * Construct an uninitialized WSLoginModuleImpl object.
     */
    public sampleSpnegoMappingLoginModule() {
        debugOut("sampleSpnegoMappingLoginModule() entry");
        debugOut("sampleSpnegoMappingLoginModule() exit");
    }

    /**
     * Initialize this login module.
     *
     * This is called by the LoginContext after this login module is
     * instantiated. The relevant information is passed from the LoginContext
     * to this login module. If the login module does not understand any of the data
     * stored in the sharedState and options parameters,
     * they can be ignored.
     *
     * @param subject The subject to be authenticated.
     * @param callbackHandler A CallbackHandler for communicating with the end user to gather login information
     * (e.g., username and password).
     * @param sharedState The state shared with other configured login modules.
     */
}
```



```

    * @param options The options specified in the login configuration for this particular login module.
    */
public void initialize(Subject subject, CallbackHandler callbackHandler,
    Map sharedState, Map options) {
    debugOut("initialize(subject = \"" + subject.toString() +
        "\", callbackHandler = \"" + callbackHandler.toString() +
        "\", sharedState = \"" + sharedState.toString() +
        "\", options = \"" + options.toString() + "\")");

    this.subject = subject;
    this.callbackHandler = callbackHandler;
    this.sharedState = sharedState;
    this.options = options;

    debug = "true".equalsIgnoreCase((String)this.options.get("debug"));

    debugOut("initialize() exit");
}

/**
 *
 * Method to authenticate a Subject (phase 1).
 *
 *
 * This method authenticates a Subject. It uses CallbackHandler to gather
 * the Subject information, like username and password for example, and verify these
 * information. The result of the authentication is saved in the private state within
 * this login module.
 *
 *
 * @return true if the authentication succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the authentication fails.
 */
public boolean login() throws LoginException
{
    debugOut("sampleSpnegoMappingLoginModule.login() entry");

    boolean succeeded = false;
    java.util.Set krb5Principals= subject.getPrincipals(KerberosPrincipal.class);
    java.util.Iterator krb5PrincIter = krb5Principals.iterator();

    while (krb5PrincIter.hasNext()) {
        Object princObj = krb5PrincIter.next();
        debugOut("Kerberos principal name: "+ princObj.toString());

        if (princObj != null && princObj.toString().equals("utle@WSSEC.AUSTIN.IBM.COM")){
            mapUid = "user1";
            debugOut("mapUid: "+mapUid);

            java.util.Hashtable customProperties = (java.util.Hashtable)
            sharedState.get(AttributeNameConstants.WSCREDENTIAL_PROPERTIES_KEY);
            if (customProperties == null) {
                customProperties = new java.util.Hashtable();
            }
            succeeded = true;
            customProperties.put(AttributeNameConstants.WSCREDENTIAL_USERID, mapUid);

            Map<String,java.util.Hashtable>mySharedState=(Map<String,java.
            util.Hashtable>)sharedState;
            mySharedState.put(AttributeNameConstants.WSCREDENTIAL_PROPERTIES_KEY,
            customProperties);

            debugOut("Add a mapping user ID to Hashtable, mapping ID = "+mapUid);

```

```

        debugOut("login() custom properties = " + customProperties);
    }
}

succeeded = true;
debugOut("sampleSpnegoMappingLoginModule.login() exit");

return succeeded;
}

/**
 *
 * Method to commit the authentication result (phase 2).
 *
 *
 * This method is called if the LoginContext's overall authentication
 * succeeded (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL login module
 * succeeded).
 *
 * @return true if the commit succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the commit fails.
 */
public boolean commit() throws LoginException
{
    debugOut("commit()");

    debugOut("commit()");

    return true;
}

/**
 * Method to abort the authentication process (phase 2).
 *
 *
 * This method is called if the LoginContext's overall authentication
 * failed (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL login module
 * did not succeed).
 *
 *
 * If this login module's authentication attempt succeeded, then this method cleans
 * up the previous state saved in phase 1.
 *
 * @return true if the abort succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the abort fails.
 */
public boolean abort() throws LoginException {
    debugOut("abort() entry");
    debugOut("abort() exit");
    return true;
}

/**
 * Method which logs out a Subject.
 *
 * @return true if the logout succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the logout fails.

```

```

    */
    public boolean logout() throws LoginException
    {
        debugOut("logout() entry");
        debugOut("logout() exit");

        return true;
    }

    private void cleanup()
    {
        debugOut("cleanup() entry");
        debugOut("cleanup() exit");
    }

    /*
     *
     * Private method to print trace information. This implementation uses System.out
     * to print trace information to standard output, but a custom tracing system can
     * be implemented here as well.
     *
     */
    private void debugOut(Object o)
    {
        System.out.println("Debug: " + MAPPING_MODULE_NAME);
        if (o != null) {
            if (o.getClass().isArray()) {
                int length = Array.getLength(o);
                for (int i = 0; i < length; i++) {
                    System.out.println("\t" + Array.get(o, i));
                }
            } else {
                System.out.println("\t" + o);
            }
        }
    }

    private Subject subject;
    private CallbackHandler callbackHandler;
    private Map sharedState;
    private Map options;

    protected boolean debug = false;
}

```

2. For SPNEGO Web authentication without Kerberos authentication, this JAAS custom login module must first be inserted in the stack login for system login config for WEB_INBOUND, RMI_INBOUND and DEFAULT.
3. For Kerberos authentication, this JAAS custom login module must be inserted immediately before the ItpaLoginModule for system login config for WEB_INBOUND, RMI_INBOUND and DEFAULT.

Results

Using the custom login module, the Kerberos principal name is mapped to the WebSphere Application Server's security registry.

Configuring a Java client for Kerberos authentication

A Java client can authenticate with WebSphere Application server with a Kerberos principal name and password or with the Kerberos credential cache (krb5Ccache).

1. Create a Kerberos configuration file (**krb5.ini** or **krb5.conf**). Read about "Creating a Kerberos configuration file" on page 270 for more information.
2. Place either the **krb5.ini** or **krb5.conf** files you have created in a default location. If either file is not located in the default location you must set `com.ibm.COBRA.krb5ConfigFile` in the `sas.client.props` file with the correct path and Kerberos configuration file name.

On a Windows operating system, the default location is `c:\winnt\krb5.ini`.

On a Linux operating system, the default location is `/etc/krb5.conf`.

On other UNIX-based operating systems, the default location is `/etc/krb5/krb5.conf`.

On the z/OS operating system, the default location is `/etc/krb5/krb5.conf`.

On the i5/OS operating system, the default location is `/QIBM/UserData/OS400/NetworkAuthentication/krb5.conf`

3. In the **sas.client.props** file, set the `com.ibm.CORBA.authenticationTarget` property to `KRB5`. Read about Chapter 10, “Configuring security with scripting,” on page 875 for more information.
4. Also in the **sas.client.props** file, set the `com.ibm.CORBA.loginSource` property to one of the supported values shown below:

When `authenticationTarget` is `BasicAuth`, the `loginSource` supported are:

- `prompt` [default]
- `properties`
- `stdin`
- `none`

When `authenticationTarget` is `KRB5`, the `loginSource` supported are:

- `prompt` [default]
- `properties`
- `stdin`
- `none`
- `krb5Ccache`
- `krb5Ccache:prompt`
- `krb5Ccache:properties`
- `krb5Ccache:stdin`

Also consider the following:

krb5Ccache:prompt

Use `krb5Ccache` to authenticate to WebSphere Application Server first. If it fails, then it falls back to `prompt`.

krb5Ccache:properties

Use `krb5Ccache` to authenticate to WebSphere Application Server first. If it fails, then it falls back to `properties`.

krb5Ccache:stdin

Use `krb5Ccache` to authenticate to WebSphere Application Server first. If it fails, then it falls back to `stdin`.

5. If the `authenticationTarget` is `KRB5` and `loginSource` is the Kerberos credential cache, do the following:
 - a. In the **wsjaas_client.conf** file, update the `WSKRB5Login` entry:

```
WSKRB5Login{
    com.ibm.ws.security.auth.kerberos.Krb5LoginModuleWrapperClient required
    credsType=INITIATOR useFirstPass=false
    forwardable=false renewable=false noAddress=false;
};
```

- b. If the Kerberos credential cache is not at the default location, set the `com.ibm.CORBA.krb5CcacheFile` property as a URL: For example:

```
com.ibm.CORBA.krb5CcacheFile=FILE:/home/smith/krb5cc_smith
```

The default location of the Kerberos credential cache file depends on which operating system you use. The user credential cache is located in the following order:

The file referenced by the Java property KRB5CCNAME
<user.home>/krb5cc_<user.name>
<user.home>/krb5cc (if <user.name> cannot be obtained)

- Optional: If an error occurs, make sure that all WebSphere Application Server machine's clocks are in sync with the KDC machines. Validate a Kerberos principal name and password by using the **kinit** command. It is recommended that you use the **kinit** command that comes with the Java SE Development Kit (JDK) 6.

Type the following to receive help for this command:

```
kinit -help
```

An example of the **kinit** command is as follows. In this example, a Kerberos ticket-granting ticket (TGT) can be obtained for duke, and the Kerberos credential cache is stored in the default location, c:\Documents and Settings\duke\krb5cc_duke:

```
kinit duke@JAVA.SUN.COM
```

Results

You have now configured a Java client for Kerberos authentication.

You can complete your configuration of Kerberos on the server side by using either the administrative console or by using wsadmin commands. Read about "Configuring Kerberos as the authentication mechanism using the administrative console" on page 266 or "Kerberos authentication commands" on page 1262 respectively for more information.

RSA token authentication mechanism

An *authentication mechanism* defines rules about security information, such as whether a credential is forwardable to another Java process, and the format of how security information is stored in both credentials and tokens. The Rivest Shamir Adleman (RSA) token authentication mechanism simplifies the security environment for flexible management topology, that is, the topology where you can locally or remotely submit and manage administrative jobs through a job manager that manages applications, perform product maintenance, modify configurations, and control the application server runtime. The RSA authentication mechanism is only used for server-to-server administrative authentication (admin connector and file transfer requests). The RSA authentication mechanism does not replace LTPA or Kerberos for use by applications.

Note: The RSA token authentication mechanism is new to this release of WebSphere Application Server. It aids the flexible management objective to preserve the base profiles configurations and isolate them from a security perspective. This mechanism permits the base profiles managed by an administrative agent to have different Lightweight Third-Party Authentication (LTPA) keys, different user registries, and different administrative users.

Authentication is the process of establishing whether a client is who or what it claims to be in a particular context. A client can be either an end user, a machine, or an application. An authentication mechanism in WebSphere Application Server typically collaborates closely with a *user registry*. The user registry is the user and groups account repository that the authentication mechanism consults with when performing authentication. The authentication mechanism is responsible for creating a *credential*, which is an internal product representation of a successfully authenticated client user. Not all credentials are created equally. The abilities of the credential are determined by the configured authentication mechanism.

Authentication process

The RSA token authentication mechanism ensures that after the RSA root signer certificate (20 year lifetime) is exchanged between two administrative processes, there is no need to synchronize security information among disparate profiles for administrative requests. The RSA personal certificate (1 year lifetime) is used to perform the cryptographic operations on the RSA tokens and can be verified by the

long-lived RSA root. RSA token authentication is different from LTPA where keys are shared and if one side changes, all sides need to change. Since RSA token authentication is based on a PKI infrastructure, it benefits from the scalability and manageability of this technology in a large topology.

An RSA token has more advanced security features than LTPA; this includes a nonce value that makes it a one-time use token, a short expiration period (since it's a one-time use token), and trust, which is established based on certificates in the target RSA trust store.

RSA token authentication does not use the same certificates as used by Secure Sockets Layer (SSL). This is the reason RSA has its own keystores. In order to isolate the trust established for RSA, the trust store, keystore, and root keystore, need to be different from the SSL configuration.

Note: SSL personal certificates given to pure clients are often signed by the same SSL root certificate used by servers, and this allows a pure client to send an RSA token to a server and act as an administrator. This should be avoided for the RSA token authentication mechanism. The RSA token authentication mechanism has its own root certificate which signs personal certificates that are used to encrypt and sign parts of the token.

The data stored in an RSA token is based on the identity of the client subject. The client subject can be based on LTPA or Kerberos, but the RSA token does not use this protection for administrative requests. The RSA token is easier to use while still maintaining a secure transportation of the identity. The data in an RSA token includes:

- Version
- Nonce
- Expiration
- Realm
- Principal
- Access ID
- Roles (not currently used)
- Groups
- Custom data

Custom data can be added to the `WSCredential` on the sending side (prior to going outbound) by creating a properties object, adding custom attributes, and adding this to the `WSCredential` in the following way.

```
import com.ibm.websphere.security.cred.WSCredential;

java.util.Properties props = new java.util.Properties();
props.setProperty("myAttribute", "myValue");
WSCredential.put ("customRSAProperties", props);
```

Once the Subject is created at the target process, you can get access to these attributes in the following way.

```
java.util.Properties props = (java.util.Properties) WSCredential.get("customRSAProperties");
```

This data is placed into a hash table at the target side and the hash table is used in a Java™ Authentication and Authorization Service (JAAS) login to obtain a subject at the target that contains the same attributes from the RSA token. With the target containing the same attributes from the RSA token, you can have a subject at the target side that is not from the same realm used by the target. For this authorization to succeed, a cross-realm mapping is required within the administrative authorization table unless the identity is a trusted server ID.

The figure (below) is an overview of the RSA token authentication mechanism and describes the process that takes place when a request is sent from a server-as-client to a target server. The server-as-client has an administrative subject on the thread that is used as input to create the RSA token. The other

information needed is RSA public certificate of the target server. This certificate must be retrieved by making a “bootstrap” MBean request to the target process prior to sending any real requests. The target bootstrap request retrieves the public certificate from the target process. When creating an RSA token, the primary purpose of obtaining the target’s public certificate is to encrypt the secret key. Only the target can decrypt the secret key, which is used to encrypt the user data.

The client’s private key is used to sign both the secret key and the user data. The client’s public key is embedded in the RSA token and validated at the target. If the client’s public key is not trusted when calling the CertPath APIs at the target, the RSA token validation cannot continue. If the client’s public key is trusted, it can be used to verify the secret key and user data signatures.

The basic goal is to convert the client subject into a subject at the target by securely propagating the required information. After the subject is generated at the target, the RSA authentication mechanism process is complete.

Configuring the RSA token authentication mechanism

You use the WebSphere Application Server administrative console to configure the Rivest Shamir Adleman (RSA) token authentication mechanism. The RSA token authentication mechanism can only be used for administrative requests. As such, the authentication mechanism choices for administrative authentication are part of the Global Security panel of the administrative console.

Before you begin

RSA token authentication mechanism is the default selection for the application server, administrative agent, and job manager profiles. LTPA is still the default for the deployment manager profile to preserve the same behavior for the existing topology.

About this task

You configure Lightweight Third-Party Authentication (LTPA) and Kerberos on the main authentication mechanism panels of the administrative console as well as configure RSA token authentication. During registration of a base profile with the administrative agent, the trusted certificates on both sides are updated with the root signer for the other. The same process occurs during registration of an administrative agent or deployment manager with a job manager. When removing the registration, the trusted signers are removed from both sides so that trust is no longer established.

By default, the RSA mechanism is set up correctly during the registration tasks, such as **registerNode** or **registerWithJobManager**. No further actions are necessary to establish trust within these environments. However, if you need to establish trust between two base servers or between two admin agents, for example, you can use the following steps to further configure the RSA token authentication mechanism:

1. Click **Security > Global security** . Under Administrative security click the link to **Administrative authentication**.
2. Select the RSA token radio button. Select a data encryption keystore from the drop-down list. The option is recommend for flexible systems administration.
3. Select the trusted signers keystore from the drop-down list.
4. Enter the nonce cache timeout value.
5. Enter token timeout value.
6. Click **Apply** and **Save**.

Results

You configured the use of the RSA token authentication mechanism.

RSA token authentication settings

Use this panel to configure RSA token authentication.

To view this administrative console page, click **Security > Global security**. Under Administrative security click **Administrative authentication**.

The administrative authentication method is used when an administrative process on this profile connects to another profile. If the primary authentication method is set to RSA token and that primary method fails, the system attempts to use the current application authentication method (which could be SWAM, Kerberos, or LTPA for example).

Note: SWAM is deprecated and will be removed in a future release.

RSA token (recommended for flexible systems administration):

RSA token is an authentication mechanism using certificates for signing and encryption portions of the security information being propagated.

Default: Enabled

Data encryption keystore:

This is the keystore that contains the personal certificate used to encrypt and sign RSA tokens.

Data type: text

Personal certificate for encryption:

This is the alias found in the Data encryption keystore that is used to encrypt and sign RSA tokens.

Data type: text

Trusted signers keystore:

This is the keystore used to contain signer certificates that can validate RSA tokens sent by other servers. The RSA token contains a sending certificate that needs to be validated by this trust store using a CertPath validation.

Data type: text

Nonce cache timeout:

Specifies the amount of time, in minutes, that the issued token is valid.

This field displays the maximum timeout, in minutes, for a token to be considered valid.

Data type: Integer
Default: 20
Minimum: 10
Maximum: Integer.MAX_VALUE

Token timeout:

Specifies the amount of time, in minutes, that the issued token is valid.

This field displays the maximum timeout, in minutes, for a token to be considered valid.

Data type:	Integer
Default:	10
Minimum:	10
Maximum:	Integer.MAX_VALUE

Only use the active application authentication mechanism (currently LTPA):

Select to encrypt authentication information so that the application server can send the data from one server to another in a secure manner.

The encryption of authentication information that is exchanged between servers involves the Lightweight Third-Party Authentication (LTPA) mechanism.

Kerberos:

Select to encrypt authentication information so that the application server can send the data from one server to another in a secure manner.

The encryption of authentication information that is exchanged between servers involves the Kerberos mechanism.

Note: Kerberos must be configured before this option can be selected.

RSA token certificate use

The Rivest Shamir Adleman (RSA) token uses certificates in a similar way that Secure Sockets Layer (SSL) uses them. However, the trust established for SSL and RSA are different, and RSA certificates should not use SSL certificates and vice versa. The SSL certificates can be used by pure clients, and when used for the RSA mechanism would allow the client to send an RSA token to the server. The RSA token authentication mechanism is purely for server-to-server requests and should not be used by pure clients. The way to prevent this is to control the certificates used by RSA in such a way so they are never distributed to any clients. There is a different root certificate for RSA that prevents trust being established with clients who only need SSL certificates.

RSA root certificate

For each profile there is a root certificate stored in the **rsatoken-root-key.p12** keystore. The sole purpose of this RSA root certificate is to sign the RSA personal certificate which is stored in the **rsatoken-key.p12** keystore. The RSA root certificate has a default lifetime of 20 years. The signer from the RSA root certificate is shared with other processes to establish trust.

The keytool utility is available using the QShell Interpreter. Using the keytool utility, you can list the contents of these keystores and display the keyEntry (personal certificate). The example below illustrated how this is accomplished for the **rsatoken-root-key.p12** (RSA root certificate) and **rsatoken-key.p12** (RSA personal certificate).

```
 ${profile_root}\config\cells\${cellname}\nodes\${nodename}> keytool -list -v -keystore rsatoken-root-key.p12 -storepass WebAS -storetype PKCS12
```

```
Alias name: root
Entry type: keyEntry
Certificate[1]:
```

```

Owner: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Serial number: 3474fccaf789d
Valid from: 11/12/07 2:50 PM until: 11/7/27 2:50 PM
Certificate fingerprints:
    MD5: 7E:E6:C7:E8:40:4E:9B:96:5A:66:E5:0B:37:0B:08:FD
    SHA1: 36:94:81:55:C4:48:83:27:89:C7:16:D2:AD:3D:3E:67:DF:1D:6E:87

```

```

${profile_root}\config\cells\${cellname}\nodes\${nodename}> keytool -list -v -keystore rsatoken-key.p12
-storepass WebAS -storetype PKCS12

```

```

Alias name: default
Entry type: keyEntry
Certificate[1]:
Owner: CN=9.41.62.64, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Serial number: 3475073488921
Valid from: 11/12/07 2:50 PM until: 11/11/08 2:50 PM
Certificate fingerprints:
    MD5: FF:1C:42:E3:DA:FF:DC:A4:35:B2:33:30:D1:6E:E0:19
    SHA1: A4:FB:9D:7B:A1:5B:6A:37:9F:20:BD:B2:BD:98:FA:68:71:57:28:62
Certificate[2]:
Owner: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode
04, O=IBM, C=US
Serial number: 3474fccaf789d
Valid from: 11/12/07 2:50 PM until: 11/7/27 2:50 PM
Certificate fingerprints:
    MD5: 7E:E6:C7:E8:40:4E:9B:96:5A:66:E5:0B:37:0B:08:FD
    SHA1: 36:94:81:55:C4:48:83:27:89:C7:16:D2:AD:3D:3E:67:DF:1D:6E:87

```

The purpose of the RSA personal certificate is to sign and encrypt information in the RSA token. The RSA personal certificate has a default lifetime of one year because it is used to sign and encrypt data that is transmitted over the wire. Refreshing the certificate is performed by the certificate expiration monitor, which is used for any other certificate in the system including SSL certificates.

RSA token trust is established when the **rsatoken-trust.p12** of the target process contains the signer of the root certificate of the client process that sends a token. Inside the RSA token is the public certificate of the client, which must be validated at the target before being used to decrypt data. The validation of the client's public certificate is performed using the CertPath APIs, which use the **rsatoken-trust.p12** as the source of certificates used during the validation.

The following example shows the use of the keytool utility to list the **rsatoken-trust.p12** keystore.

Note: This trust store contains three trustedCertEntry (public certificate) entries. The root public certificate from the administrative agent, a root public certificate from a job manager to which it is registered, and a root public certificate from a base profile to which it is registered.

```

${profile_root}\config\cells\${cellname}\nodes\${nodename}> keytool -list -v -keystore rsatoken-trust.p12
-storepass WebAS -storetype PKCS12

```

```

Alias name: root
Entry type: trustedCertEntry

Owner: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60AACell04, OU=BIRKT60AANode04, O=IBM, C=US
Serial number: 3474fccaf789d
Valid from: 11/12/07 2:50 PM until: 11/7/27 2:50 PM
Certificate fingerprints:
    MD5: 7E:E6:C7:E8:40:4E:9B:96:5A:66:E5:0B:37:0B:08:FD
    SHA1: 36:94:81:55:C4:48:83:27:89:C7:16:D2:AD:3D:3E:67:DF:1D:6E:87
*****
Alias name: cn=9.41.62.64, ou=root certificate, ou=birkt60jobmgrcell02, ou=birkt
60jobmgr02, o=ibm, c=us

```

Entry type: trustedCertEntry

Owner: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60JobMgrCell02, OU=BIRKT60JobMgr02, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60JobMgrCell02, OU=BIRKT60JobMgr02, O=IBM, C=US
Serial number: 34cc4c5d71740
Valid from: 11/12/07 4:30 PM until: 11/7/27 4:30 PM
Certificate fingerprints:
MD5: AB:65:3A:04:5B:C7:6D:A8:B1:98:B9:7B:65:A8:FA:F8
SHA1: C0:83:FE:D0:B6:30:FB:A1:10:41:4B:8E:50:4B:78:40:0F:E5:E3:35

Alias name: birk60node19_signer

Entry type: trustedCertEntry

Owner: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60Node15Cell, OU=BIRKT60Node19, O=IBM, C=US
Issuer: CN=9.41.62.64, OU=Root Certificate, OU=BIRKT60Node15Cell, OU=BIRKT60Node19, O=IBM, C=US
Serial number: 34825d997fda3
Valid from: 11/12/07 3:06 PM until: 11/7/27 3:06 PM
Certificate fingerprints:
MD5: 66:61:CE:7C:C7:44:8B:A7:23:FF:1B:68:E4:AC:24:55
SHA1: 25:E0:6B:D9:60:BB:67:5B:C6:67:BD:02:2C:54:E3:DA:24:E5:31:A3

You can replace the RSA personal certificate in the rsa-key.p12 and the public key in the rsa-trust.p12 with your own personal certificate. If you do this prior to federation to an administrative agent or job manager, the exchange of certificates is done for you. If you change the certificate after federation, you need to make sure the rsa-trust.p12 on the administrative agent or job manager is updated with the signer for your new certificate to establish trust.

Message layer authentication

Defines the credential information and sends that information across the network so that a receiving server can interpret it.

When you send authentication information across the network using a token the transmission is considered message layer authentication because the data is sent with the message inside a service context.

A pure Java client uses Kerberos (KRB5) or basic authentication, or Generic Security Services Username Password (GSSUP), as the authentication mechanism to establish client identity.

However, a servlet can use either basic authentication (GSSUP) or the authentication mechanism of the server, Kerberos (KRB5) or Lightweight Third Party Authentication (LTPA), to send security information in the message layer. Use KRB5 or LTPA by authenticating or by mapping the basic authentication credentials to the security mechanism of the server.

The security token that is contained in a token-based credential is authentication mechanism-specific. The way that the token is interpreted is only known by the authentication mechanism. Therefore, each authentication mechanism has an object ID (OID) representing it. The OID and the client token are sent to the server, so that the server knows which mechanism to use when reading and validating the token. The following list contains the OIDs for each mechanism:

BasicAuth (GSSUP): oid:2.23.130.1.1.1

KRB5: OID: 1.2.840.113554.1.2.2

LTPA: oid:1.3.18.0.2.30.2

SWAM: No OID because it is not forwardable

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

On the server, the authentication mechanisms can interpret the token and create a credential, or they can authenticate basic authentication data from the client, and create a credential. Either way, the created credential is the *received* credential that the authorization check uses to determine if the user has access to invoke the method. You can specify the authentication mechanism by using the following property on the client side:

- `com.ibm.CORBA.authenticationTarget`

Basic authentication (BasicAuth) and KRB5 are currently the only valid values. You can configure the server through the administrative console.

Note: When **perform basic authentication** is enabled, if the client is not similarly configured (and does not pass a credential such as a user ID and password), the server object request broker (ORB) does not.

Configuring authentication retries

Situations occur where you want a prompt to display again if you entered your user ID and password incorrectly or you want a method to retry when a particular error occurs back at the client. If you can correct the error by information at the client side, the system automatically performs a retry without the client seeing the failure, if the system is configured appropriately.

Some of these errors include:

- Entering a user ID and password that are not valid
- Having an expired credential on the server
- Failing to find the stateful session on the server

By default, authentication retries are enabled and perform three retries before returning the error to the client. Use the `com.ibm.CORBA.authenticationRetryEnabled` property (True or False) to enable or disable authentication retries. Use the `com.ibm.CORBA.authenticationRetryCount` property to specify the number of retry attempts.

Immediate validating of a basic authentication login

In WebSphere Application Server Version 6.x, a behavior is defined during `request_login` for a BasicAuth login. In releases prior to Version 5, a BasicAuth login takes the user ID and password entered through the `loginSource` method and creates a BasicAuth credential. If either the user ID or the password is not valid, the client program does not find out until the first method request is attempted. When the user ID or password is specified during a prompt or programmatic login, the user ID and password are authenticated by default with the security server, with a True or False returned as the result. If False, an `org.omg.SecurityLevel2.LoginFailed` exception is returned to the client indicating that the user ID and password are not valid. If True, then the BasicAuth credential is returned to the caller of the `request_login`. To disable this feature on the pure client, specify `com.ibm.CORBA.validateBasicAuth=false`. By default, this feature is set to True. On the server side, specify this property in the security dynamic properties.

Note: Set `com.ibm.CORBA.validateBasicAuth=false` whenever connecting to a z/OS server. This function does not currently work from a distributed client to a z/OS server since the SecurityServer is located using the "UNAUTHENTICATED" principal, which is not accepted on a z/OS system.

Trust associations

Trust association enables the integration of IBM WebSphere Application Server security and third-party security servers. More specifically, a reverse proxy server can act as a front-end authentication server while the product applies its own authorization policy onto the resulting credentials that are passed by the proxy server.

Demand for such an integrated configuration has become more compelling, especially when a single product cannot meet all of the customer needs or when migration is not a viable solution. This article provides a conceptual background behind the approach.

In this setup, WebSphere Application Server is used as a back-end server to further exploit its fine-grained access control. The reverse proxy server passes the HTTP request to WebSphere Application Server that includes the credentials of the authenticated user. WebSphere Application Server then uses these credentials to authorize the request.

Trust association model

The idea that WebSphere Application Server can support trust association implies that the product application security recognizes and processes HTTP requests that are received from a reverse proxy server. WebSphere Application Server and the proxy server engage in a contract in which the product gives its full trust to the proxy server and the proxy server applies its authentication policies on every Web request that is dispatched to WebSphere Application Server. This trust is validated by the interceptors that reside in the product environment for every request received. The method of validation is agreed upon by the proxy server and the interceptor.

Running in trust association mode does not prohibit WebSphere Application Server from accepting requests that did not pass through the proxy server. In this case, no interceptor is needed for validating trust.

WebSphere Application Server supports the following trust association interceptor (TAI) interfaces:

com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus

This TAI interceptor implementation that implements the new WebSphere Application Server interface supports WebSphere Application Server Version 5.1.1 and later. The interface supports WebSEAL Version 5.1, but does not support WebSEAL Version 4.1. For an explanation of security attribute propagation, see “Security attribute propagation” on page 448.

com.ibm.ws.security.spnego.TrustAssociationInterceptorImpl

This interceptor is new to this release. SPNEGO has replaced SPNEGO TAI as the Web authenticator for WebSphere Application Server.

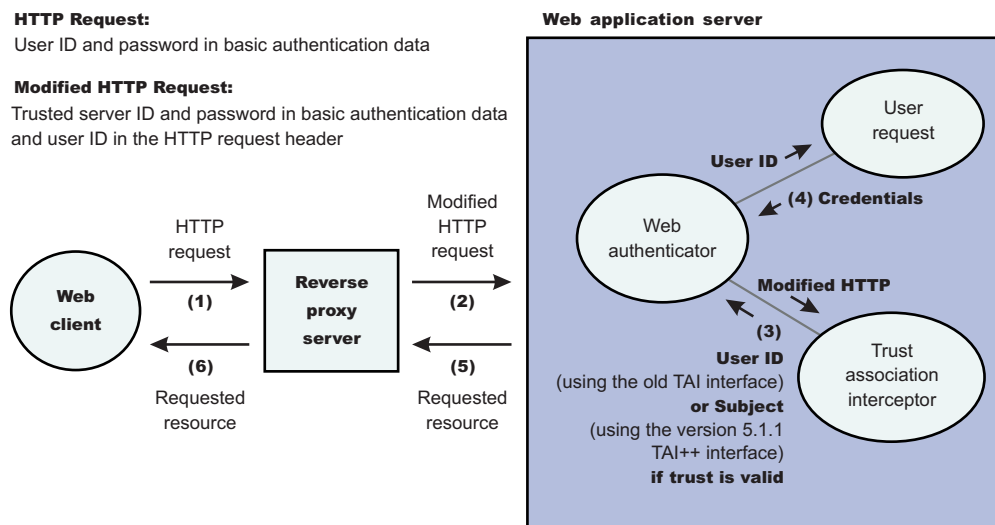
Trust association model

HTTP Request:

User ID and password in basic authentication data

Modified HTTP Request:

Trusted server ID and password in basic authentication data and user ID in the HTTP request header



IBM WebSphere Application Server: WebSEAL Integration

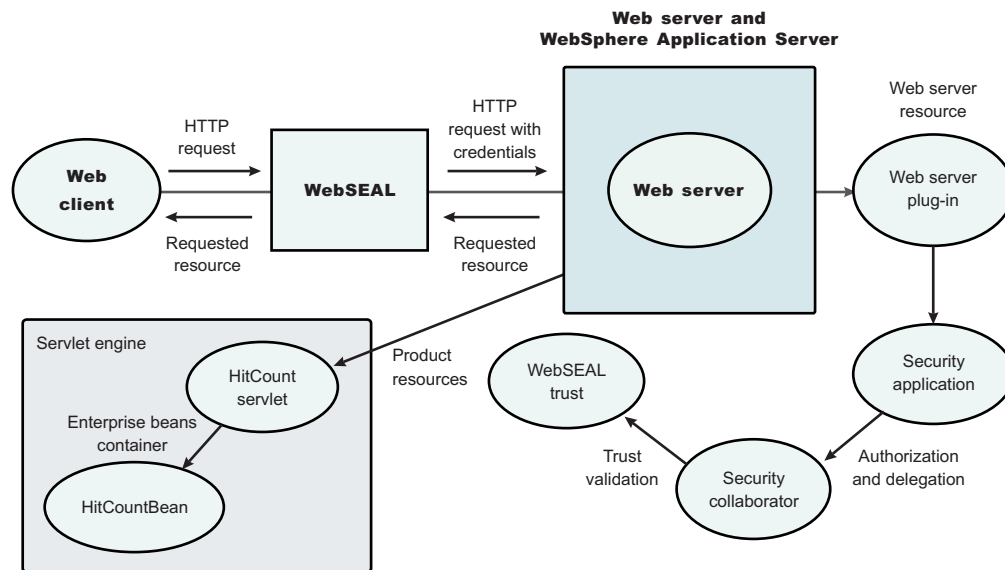
The integration of WebSEAL and WebSphere Application Server security is achieved by placing the WebSEAL server at the front-end as a reverse proxy server. From a WebSEAL management perspective, a junction is created with WebSEAL on one end, and the product Web server on the other end. A junction is a logical connection that is created to establish a path from the WebSEAL server to another server.

In this setup, a request for Web resources that are stored in a protected domain of the product is submitted to the WebSEAL server where it is authenticated against the WebSEAL security realm. If the requesting user has access to the junction, the request is transmitted to the WebSphere Application Server HTTP server through the junction, and then to the application server.

Meanwhile, WebSphere Application Server validates every request that comes through the junction to ensure that the source is a trusted party. This process is referenced as *validating the trust* and it is performed by a WebSEAL product-designated interceptor. If the validation is successful, WebSphere Application Server authorizes the request by checking whether the client user has the required permissions to access the Web resource. If so, the Web resource is delivered to the WebSEAL server through the Web server, which then gives the resource to the client user.

WebSEAL server


The policy director delegates all of the Web requests to its Web component, the WebSEAL server. One of the major functions of the server is to perform authentication of the requesting user. The WebSEAL server consults a Lightweight Directory Access Protocol (LDAP) directory. It can also map the original user ID to another user ID, such as when global single sign-on (GSO) is used.



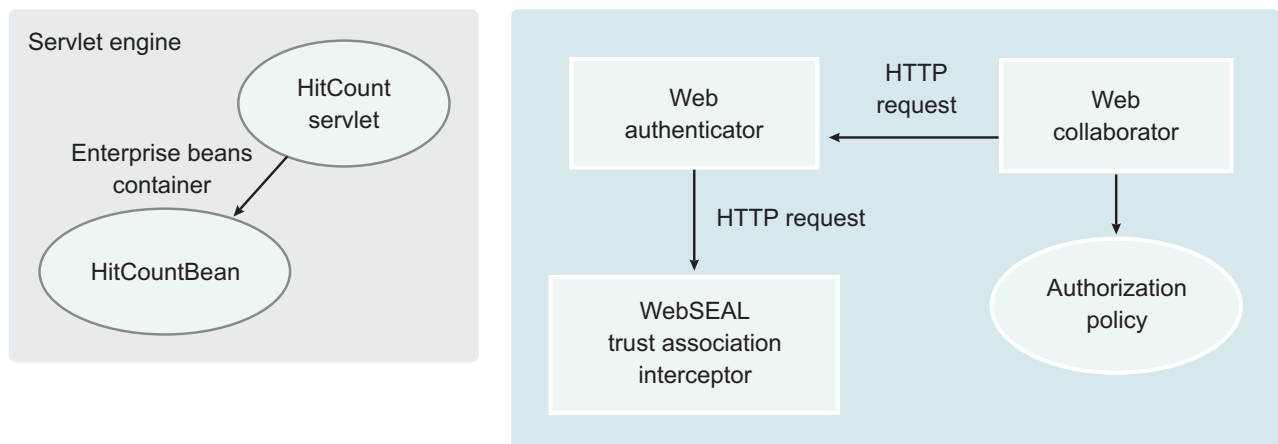
For successful authentication, the server plays the role of a client to WebSphere Application Server when channeling the request. The server needs its own user ID and password to identify itself to WebSphere Application Server. This identity must be valid in the security realm of WebSphere Application Server. The WebSEAL server replaces the basic authentication information in the HTTP request with its own user ID and password. In addition, WebSphere Application Server must determine the credentials of the requesting client so that the application server has an identity to use as a basis for its authorization decisions. This information is transmitted through the HTTP request by creating a header called `iv-creds`, with the Tivoli Access Manager user credentials as its value.

HTTP server

The junction that is created in the WebSEAL server must get to the HTTP server that serves as the product front end. However, the HTTP server is shielded from knowing that trust association is used. As far as it is concerned, the WebSEAL product is just another HTTP client, and as part of its normal routines, it sends the HTTP request to the product. The only requirement on the HTTP server is a Secure Sockets Layer (SSL) configuration using server authentication only. This requirement protects the requests that flow within the junction.



WebSphere Application Server



Web collaborator

When trust association is enabled, the Web collaborator manages the interceptors that are configured in the system. The Web collaborator loads and initializes these interceptors when you restart your servers. When a request is passed to WebSphere Application Server by the Web server, the Web collaborator eventually receives the request for a security check. Two actions must take place:

1. The request must be authenticated.
2. The request must be authorized.

The Web authenticator is called to authenticate the request by passing the HTTP request. If successful, a good credential record is returned by the authenticator, which the Web collaborator uses to base its authorization for the requested resource. If the authorization succeeds, the Web collaborator indicates to WebSphere Application Server that the security check has succeeded and that the requested resource can be served.

Web authenticator

The Web authenticator is asked by the Web collaborator to authenticate a given HTTP request. Knowing that trust association is enabled, the task of the Web authenticator is to find the appropriate trust association interceptor to direct the request for processing. The Web authenticator queries every available interceptor. If no target interceptor is found, the Web authenticator processes the request as though trust association is not enabled.

Note:

WebSphere Application Server Version 4 through WebSphere Application Server Version 6.x support the `com.ibm.websphere.security.TrustAssociationInterceptor.java` interface. WebSphere Application Server Version 7.0.x and later supports the `com.ibm.ws.security.spnego.TrustAssociationInterceptorImpl` interface.

Trust association interceptor interface

The intent of the trust association interceptor interface is to have reverse proxy security servers (RPSS) exist as the exposed entry points to perform authentication and coarse-grained authorization, while WebSphere Application Server enforces further fine-grained access control. Trust associations improve security by reducing the scope and risk of exposure.

In a typical e-business infrastructure, the distributed environment of a company consists of Web application servers, Web servers, existing systems, and one or more RPSS, such as the Tivoli WebSEAL product. Such reverse proxy servers, front-end security servers, or security plug-ins registered within Web servers, guard the HTTP access requests to the Web servers and the Web application servers. While protecting access to the Uniform Resource Identifiers (URIs), these RPSS perform authentication, coarse-grained authorization, and request routing to the target application server.

Integrating third-party HTTP reverse proxy servers

These steps are required to use a trust association interceptor with a reverse proxy security server.

About this task

WebSphere Application Server enables you to use multiple trust association interceptors. The application server uses the first interceptor that can handle the request.

1. Access the administrative console.
Type `http://fully_qualified_host_name:port_number/ibm/console` in a Web browser.
Port 9060 is the default port number for accessing the administrative console. During installation, however, you might have specified a different port number. Use the appropriate port number.
2. Click **Security > Global security**.
3. Under Web and SIP security, click **Trust association**.
4. Select the **Enable trust association** option.
5. Under Additional properties, click **Interceptors**. The default value appears.
6. Verify that the appropriate trust association interceptors are listed.

Results

Trust association is enabled.

What to do next

1. If you are enabling security, make sure that you complete the remaining steps for enabling security.
2. Save, stop and restart all of the product servers (deployment managers, nodes and application servers) for the changes to take effect.

Trust association settings

Use this page to enable trust association, which integrates application server security and third-party security servers. More specifically, a reverse proxy server can act as a front-end authentication server while the product applies its own authorization policy onto the resulting credentials passed by the proxy server.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand Web security and click **Trust association**.

When security is enabled and any of these properties change, go to the Global security panel and click **Apply** to validate the changes.

Enable trust association

Specifies whether trust association is enabled.

Data type:	Boolean
Default:	Disable
Range:	Enable or Disable

Trust association interceptor collection

Use this page to specify trust information for reverse security proxy servers.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand Web and SIP security and click **Trust association**.
3. Under Additional Properties, click **Interceptors**.

When security is enabled and any of these properties are changed, go to the Global security panel and click **Apply** to validate the changes.

Interceptor class name

Specifies the trust association interceptor class name.

Data type
String

Trust association interceptor settings

Use this page to specify trust information for reverse security proxy servers.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Web and SIP security**.
3. Click **Trust association**.
4. Under Additional Properties, click **Interceptors > New**.

Interceptor class name

Specifies the trust association interceptor class name.

Data type
String

Single sign-on

With single sign-on (SSO) support, Web users can authenticate once when accessing both WebSphere Application Server resources, such as HTML, JavaServer Pages (JSP) files, servlets, enterprise beans, and Lotus Domino resources, such as documents in a Domino database, or accessing resources in multiple WebSphere Application Server domains.

There are various ways to accomplish SSO, with the most common in WebSphere using LTPA cookies. LTPA cookies do not require any particular client and allow SSO across different cells provide the registry and LTPA keys are the same.

There are other flavors of SSO, including Simple and Protected GSS-API Negotiation (SPNEGO), which is a way to use the token from a Kerberos login (typically Windows) to authenticate to WebSphere Application Server. This prevents the user from having to type in their userid and passwords again.

Note: In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

TAIs are also a form of single sign-on when used in combination with a Proxy server that does the front-end authentication. The TAI allows the credentials to flow to WebSphere from the Proxy server and to be used to login without the need to re-authenticate the user.

Single sign-on using LTPA cookies

With single sign-on (SSO) support, Web users can authenticate once when accessing both WebSphere Application Server resources, such as HTML, JavaServer Pages (JSP) files, servlets, enterprise beans, and Lotus Domino resources, such as documents in a Domino database, or accessing resources in multiple WebSphere Application Server domains.

Application servers distributed in multiple nodes and cells can securely communicate using the Lightweight Third Party Authentication (LTPA) protocol. LTPA is intended for distributed, multiple application server and machine environments. LTPA can support security in a distributed environment through cryptography. This support permits LTPA to encrypt, digitally sign, and securely transmit authentication-related data, and later decrypt and verify the signature.

LTPA also provides the SSO feature wherein a user is required to authenticate only once in a domain name system (DNS) domain and can access resources in other WebSphere Application Server cells without getting prompted. Web users can authenticate once to a WebSphere Application Server or to a Domino server. This authentication is accomplished by configuring WebSphere Application Servers and the Domino servers to share authentication information.

Without logging in again, Web users can access other WebSphere Application Servers or Domino servers in the same DNS domain that are enabled for SSO. You can enable SSO among WebSphere Application Servers by configuring SSO for WebSphere Application Server. To enable SSO between WebSphere Application Servers and Domino servers, you must configure SSO for both WebSphere Application Server and for Domino.

Prerequisites and conditions

To take advantage of support for SSO between WebSphere Application Servers or between WebSphere Application Server and a Domino server, applications must meet the following prerequisites and conditions:

- Verify that all servers are configured as part of the same DNS domain. The realm names on each system in the DNS domain are case sensitive and must match identically. For example, if the DNS domain is specified as `mycompany.com`, then SSO is effective with any Domino server or WebSphere Application Server on a host that is part of the `mycompany.com` domain, for example, `a.mycompany.com` and `b.mycompany.com`.
- Verify that all servers share the same registry.

This registry can be either a supported Lightweight Directory Access Protocol (LDAP) directory server or, if SSO is configured between two WebSphere Application Servers, a standalone custom

registry. Domino servers do not support standalone custom registries, but you can use a Domino-supported registry as a standalone custom registry within WebSphere Application Server.

You can use a Domino directory that is configured for LDAP access or other LDAP directories for the registry. The LDAP directory product must have WebSphere Application Server support. Supported products include both Domino and LDAP servers, such as IBM Tivoli Directory Server. Regardless of the choice to use an LDAP or a standalone custom registry, the SSO configuration is the same. The difference is in the configuration of the registry.

- Define all users in a single LDAP directory. Using multiple Domino directory assistance documents to access multiple directories also is not supported.
- Enable HTTP cookies in browsers because the authentication information that is generated by the server is transported to the browser in a cookie. The cookie is used to propagate the authentication information for the user to other servers, exempting the user from entering the authentication information for every request to a different server.
- For a Domino server:
 - Domino Release 6.5.4 for iSeries and other platforms are supported.
 - A Lotus Notes® client Release 5.0.5 or later is required for configuring the Domino server for SSO.
 - You can share authentication information across multiple Domino domains.
- For WebSphere Application Server:
 - WebSphere Application Server Version 3.5 or later for all platforms are supported.
 - You can use any HTTP Web server that is supported by WebSphere Application Server.
 - You can share authentication information across multiple product administrative domains.
 - Basic authentication (user ID and password) using the basic and form-login mechanisms is supported.
 - By default, WebSphere Application Server does a case-sensitive comparison for authorization. This comparison implies that a user who is authenticated by Domino matches the entry exactly (including the base distinguished name) in the WebSphere Application Server authorization table. If case sensitivity is not considered for the authorization, enable the **Ignore Case** property in the LDAP user registry settings.

Global single sign-on principal mapping

You can use the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager to manage authentication to enterprise information systems (EIS) such as databases, transaction processing systems, and message queue systems that are located within the WebSphere Application Server security domain. Such authentication is achieved using the global single sign-on (GSO) principal mapper Java Authentication and Authorization Service (JAAS) login module for Java Platform, Enterprise Edition (Java EE) Connector Architecture resources.

With GSO principal mapping, a special-purpose JAAS login module inserts a credential into the subject header. This credential is used by the resource adapter to authenticate to the EIS. The JAAS login module used is configured on a per-connection factory basis. The default principal mapping module retrieves the user name and password information from XML configuration files. The JACC provider for Tivoli Access Manager bypasses the credential that is stored in the Extensible Markup Language (XML) configuration files and uses the Tivoli Access Manager global sign-on (GSO) database instead to provide the authentication information for the EIS security domain.

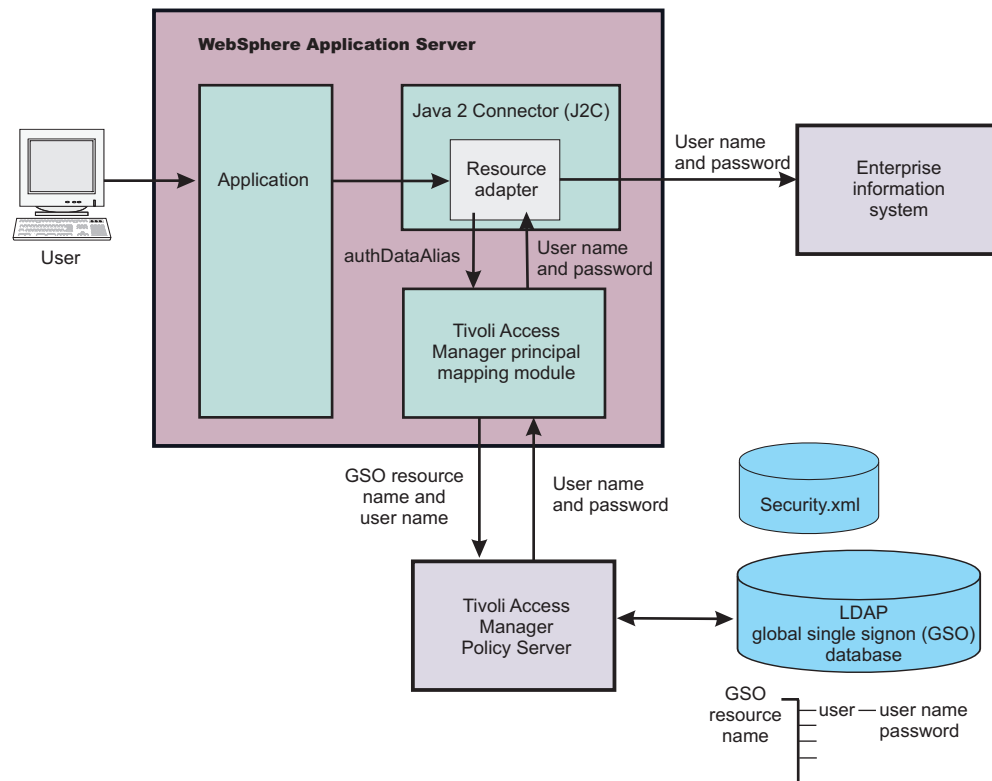
WebSphere Application Server provides a default principal mapping module that associates user credential information with EIS resources. The default mapping module is defined in the WebSphere Application Server administrative console on the Application login panel. To access the panel, click **Security > Global security**. Under Java Authentication and Authorization Service, click **Application logins**. The mapping module name is DefaultPrincipalMapping.

The EIS security domain user ID and password are defined under each connection factory by an authDataAlias attribute. The authDataAlias attribute does not contain the user name and password; this attribute contains an alias that refers to a user name and password pair that is defined elsewhere.

The Tivoli Access Manager principal mapping module uses the `authDataAlias` attribute to determine the GSO resource name and the user name that is required to perform the lookup on the Tivoli Access Manager GSO database. The Tivoli Access Manager Policy Server retrieves the GSO data from the user registry.

Tivoli Access Manager stores authentication information on the Tivoli Access Manager GSO database against a resource and user name pair.

GSO principal mapping architecture



Implementing single sign-on to minimize Web user authentications

With single sign-on (SSO) support, Web users can authenticate once when accessing Web resources across multiple WebSphere Application Servers. Form login mechanisms for Web applications require that SSO is enabled. Use this topic to configure single sign-on for the first time.

Before you begin

SSO is supported only when Lightweight Third Party Authentication (LTPA) is the authentication mechanism.

When SSO is enabled, a cookie is created containing the LTPA token and inserted into the HTTP response. When the user accesses other Web resources in any other WebSphere Application Server process in the same domain name service (DNS) domain, the cookie is sent in the request. The LTPA token is then extracted from the cookie and validated. If the request is between different cells of WebSphere Application Servers, you must share the LTPA keys and the user registry between the cells for SSO to work. The realm names on each system in the SSO domain are case sensitive and must match identically.

Windows For local OS, the realm name is the domain name if a domain is in use. If a domain is not used, the realm name is the machine name.

Linux **UNIX** The realm name is the same as the host name.

For the Lightweight Directory Access Protocol (LDAP) the realm name is the host:port realm name of the LDAP server. The LTPA authentication mechanism requires that you enable SSO if any of the Web applications have form login as the authentication method.

Because single sign-on is a subset of LTPA, it is recommended that you read “Lightweight Third Party Authentication” on page 249 for more information.

When you enable security attribute propagation, the following cookie is always added to the response:

LtpaToken2

LtpaToken2 contains stronger encryption and enables you to add multiple attributes to the token. This token contains the authentication identity and additional information such as the attributes that are used for contacting the original login server and the unique cache key for looking up the Subject when considering more than just the identity in determining uniqueness.

Note: The following cookie is optionally added to the response when the **Interoperability mode** flag is enabled:

LtpaToken

LtpaToken is used for inter-operating with previous releases of WebSphere Application Server. This token contains the authentication identity attribute only.

Note: LtpaToken is generated for releases prior to WebSphere Application Server Version 5.1.1. LtpaToken2 is generated for WebSphere Application Server Version 5.1.1 and beyond.

Token type	Purpose	How to specify
LtpaToken2 only	This is the default token type. It uses the AES-CBC-PKCS5 padding encryption strength (128-bit key size). This token is stronger than the older LtpaToken used prior to WebSphere Application Server Version 6.02. This is the recommended option when interoperability with older releases is not necessary.	Disable the Interoperability mode option in the SSO configuration panel within the administrative console. To access this panel, complete the following steps: <ol style="list-style-type: none">1. Click Security > Global security.2. Under Web security, click Single sign-on (SSO).
LtpaToken and LtpaToken2	Use to interoperate with releases prior to WebSphere Application Server Version 5.1.1. The older LtpaToken cookie is present along with the new LtpaToken2 cookie. Provided the LTPA keys are correctly shared, you should be able to interoperate with any version of WebSphere using this option.	Enable the Interoperability mode option in the SSO configuration panel within the administrative console. To access this panel, complete the following steps: <ol style="list-style-type: none">1. Click Security > Global security.2. Under Web security, click Single sign-on (SSO).

About this task

The following steps are required to configure SSO for the first time.

1. Open the administrative console.

Type `http://localhost:port_number/ibm/console` to access the administrative console in a Web browser.

Port 9060 is the default port number for accessing the administrative console. During installation, however, you might have specified a different port number. Use the appropriate port number.

2. Click **Security > Global security**.
3. Under Web security, click **Single sign-on (SSO)**.
4. Click the **Enabled** option if SSO is disabled. After you click the **Enabled** option, make sure that you complete the remaining steps to enable security.
5. Click **Requires SSL** if all of the requests are expected to use HTTPS.
6. Enter the fully qualified domain names in the **Domain name** field where SSO is effective. If you specify domain names, they must be fully qualified. If the domain name is not fully qualified, WebSphere Application Server does not set a domain name value for the LtpaToken cookie and SSO is valid only for the server that created the cookie.

When you specify multiple domains, you can use the following delimiters: a semicolon (;), a space (), a comma (,), or a pipe (|). WebSphere Application Server searches the specified domains in order from left to right. Each domain is compared with the host name of the HTTP request until the first match is located. For example, if you specify `ibm.com; austin.ibm.com` and a match is found in the `ibm.com` domain first, WebSphere Application server does not continue to search for a match in the `austin.ibm.com` domain. However, if a match is not found in either the `ibm.com` or `austin.ibm.com` domains, then WebSphere Application Server does not set a domain for the LtpaToken cookie.

You can configure the Domain name field using any of the following values:

Domain name value type	Example	Purpose
Blank		The domain is not set. This causes the browser to set the domain to the request host name. The sign-on is valid on that single host only.
Single domain name	<code>austin.ibm.com</code>	If the request is to a host within the configured domain, the sign-on is valid for all hosts within that domain. Otherwise, it is valid on the request host name only.
UseDomainFromURL	<code>UseDomainFromURL</code>	If the request is to a host within the configured domain, the sign-on is valid for all hosts within that domain. Otherwise, it is valid on the request host name only.
Multiple domain names	<code>austin.ibm.com;raleigh.ibm.com</code>	The sign-on is valid for all hosts within the domain of the request host name.
Multiple domain names and UseDomainFromURL	<ul style="list-style-type: none"> • <code>austin.ibm.com;raleigh.ibm.com; UseDomainFromURL</code> 	The sign-on is valid for all hosts within the domain of the request host name.

If you specify the `UseDomainFromURL`, WebSphere Application Server sets the SSO domain name value to the domain of the host that makes the request. For example, if an HTTP request comes from `server1.raleigh.ibm.com`, WebSphere Application Server sets the SSO domain name value to `raleigh.ibm.com`.

Note: The value, `UseDomainFromURL`, is case insensitive. You can type `usedomainfromurl` to use this value.

For more information, see “Single sign-on settings” on page 355.

7. Optional: Enable the **Interoperability mode** option if you want to support SSO connections in WebSphere Application Server version 5.1.1 or later to interoperate with previous versions of the application server. This option sets the old-style LtpaToken token into the response so it can be sent to

other servers that work only with this token type. Otherwise, only the LtpaToken2 token is added to the response. If the **Web inbound security attribute propagation** option is disabled, then only the LtpaToken token is added to the response.

- Optional: Enable the **Web inbound security attribute propagation** option if you want information added during the login at a specific front-end server to propagate to other front-end servers. The SSO token does not contain any sensitive attributes, but does understand where the original login server exists in cases where it needs to contact that server to retrieve serialized information. It also contains the cache look-up value for finding the serialized information in DynaCache, if both front-end servers are configured in the same DRS replication domain. For more information, see “Security attribute propagation” on page 448.

Note: If the following statements are true, it is recommended that you disable the **Web inbound security attribute propagation** option for performance reasons:

- You do not have any specific information added to the Subject during a login that cannot be obtained at a different front-end server.
- You did not add custom attributes to the PropagationToken token using WSSecurityHelper application programming interfaces (APIs).

If you find that you are missing custom information in the Subject, re-enable the **Web inbound security attribute propagation** option to see if the information is propagated successfully to other front-end application servers.

The following two custom properties might help to improve performance when security attribute propagation is enabled:

- com.ibm.CSI.propagateFirstCallerOnly**

When this custom property is set to `true` the first caller in the propagation token that stays on the thread is logged when security attribute propagation is enabled. Without setting this property, all of the caller switches are logged, which can affect performance.

- com.ibm.CSI.disablePropagationCallerList**

When this custom property is set to `true` the ability to add a caller or host list in the propagation token is completely disabled. This function is beneficial when the caller or host list in the propagation token is not needed in the environment.

- Click **OK**.

What to do next

For the changes to take effect, save, stop, and restart all the product deployment managers, nodes, and servers.

Creating a single sign-on for HTTP requests using SPNEGO Web authentication

Creating single sign-ons for HTTP requests using the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) Web authentication for WebSphere Application Server requires the performance of several distinct, yet related functions that when completed, allow HTTP users to log in and authenticate only once at their desktop and receive automatic authentication from the WebSphere Application Server.

Before you begin

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and

authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide the following enhancements:

- You can configure and enable SPNEGO Web authentication and filters on the WebSphere Application Server server side by using the administrative console.
- Dynamic reload of SPNEGO is provided without the need to stop and restart the WebSphere Application Server server.
- Fallback to an application login method is provided if the SPNEGO Web authentication fails.

You can enable either SPNEGO TAI or SPNEGO Web Authentication but not both.

Read about “Single sign-on for HTTP requests using SPNEGO Web authentication” on page 302 for a better understanding of what SPNEGO Web authentication is and how it is supported in this version of WebSphere Application Server.

Before starting this task, complete the following checklist:

- **Windows** A Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC).
- **Windows** A Microsoft Windows 2000 or Windows 2003 domain member (client) for example, a browser or Microsoft .NET client, that supports the SPNEGO authentication mechanism, as defined in IETF RFC 2478. Microsoft Internet Explorer Version 5.5 or later and Mozilla Firefox Version 1.0 qualify as such clients.

Note: A running domain controller and at least one client machine in that domain is required. Using SPNEGO directly from the domain controller is not supported.

- The domain member has users who can log on to the domain. Specifically, you need to have a functioning Microsoft Windows 2000 or Windows 2003 active directory domain that includes:
 - Domain controller
 - Client workstation
 - Users who can login to the client workstation
- A server platform with WebSphere Application Server running and application security enabled.
- Users on the active directory must be able to access WebSphere Application Server protected resources using a native WebSphere Application Server authentication mechanism.
- The domain controller and the host of WebSphere Application Server should have the same local time.
- Ensure the clock on clients, Microsoft Active Directory and WebSphere Application Server are synchronized to within five minutes.
- Be aware that client browsers must be SPNEGO enabled, which you perform on the client application machine (with details explained in step 2 of this task).

About this task

The objective of this machine arrangement is to permit users to successfully access WebSphere Application Server resources without having to authenticate again and thus achieve Microsoft Windows desktop single sign-on capability.

Configuring the members of this environment to establish Microsoft Windows single sign-on involves specific activities that are performed on three distinct machines:

- Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC).
- A Microsoft Windows 2000 or Windows 2003 domain member (client application), such as a browser or Microsoft .NET client.
- A server platform with WebSphere Application Server running.

Perform the following steps on the indicated machines to create single sign-on for HTTP requests using SPNEGO:

- 1. Domain Controller Machine** - Configure the Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC) This configuration activity has the following steps:
 - Create a user account for the WebSphere Application Server in a Microsoft Active Directory. This account will be eventually mapped to the Kerberos service principal name (SPN).
 - On the Microsoft Active Directory machine where the Kerberos key distribution center (KDC) is active, map the user account to the Kerberos service principal name (SPN). This user account represents the WebSphere Application Server as being a Kerberos service with the KDC. Use the Microsoft **setspn** command to map the Kerberos service principal name to a Microsoft user account.
 - Create the Kerberos keytab file and make it available to WebSphere Application Server. Use the Microsoft **ktpass** tool to create the Kerberos keytab file (krb5.keytab).

Note: You make the keytab file available to WebSphere Application Server by copying the krb5.keytab file from the Domain Controller (LDAP machine) to the WebSphere Application Server machine. Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information.

Note: After you have configured your domain controller, the following operations must lead to the following results:

 - A user account is created in the Microsoft Active Directory and mapped to a Kerberos service principal name.
 - A Kerberos keytab file (krb5.keytab) is created and made available to the WebSphere Application Server. The Kerberos keytab file contains the Kerberos service principal keys WebSphere Application Server uses to authenticate the user in the Microsoft Active Directory and the Kerberos account. Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information.
- 2. WebSphere Application Server Machine** - Configure and enable the Application Server and SPNEGO using the administrative console. Read about “Enabling and configuring SPNEGO Web authentication using the administrative console” on page 307 for more information.
- 3. Client Application Machine** - Configure the client application. Client-side applications are responsible for generating the SPNEGO token. You begin this configuration process by configuring your Web browser to use SPNEGO authentication. Read about “Configuring the client browser to use SPNEGO” on page 315 for more information.

Single sign-on for HTTP requests using SPNEGO Web authentication

You can securely negotiate and authenticate HTTP requests for secured resources in WebSphere Application Server by using the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) as the Web authentication service for WebSphere Application Server.

Note: In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide the following enhancements:

- You can configure and enable SPNEGO Web authentication and filters on WebSphere Application Server by using the administrative console.
- Dynamic reload of SPNEGO is provided without the need to stop and restart WebSphere Application Server.
- Fallback to an application login method is provided if the SPNEGO Web authentication fails.
- SPNEGO can be customized at the WebSphere security domain level. Read about “Multiple security domains” on page 78 for more information.

You can enable either SPNEGO TAI or SPNEGO Web Authentication but not both.

The following sections describe SPNEGO Web authentication in more detail:

- “What is SPNEGO?”
- “The benefits of SPNEGO Web authentication” on page 304
- “SPNEGO Web authentication in a single Kerberos realm” on page 304
- “SPNEGO Web authentication in a trusted Kerberos realm” on page 305
- “Setting up SPNEGO as the Web authentication mechanism for WebSphere Application Server” on page 306

What is SPNEGO?

SPNEGO is a standard specification defined in The Simple and Protected GSS-API Negotiation Mechanism (IETF RFC 2478).

When WebSphere Application Server global and application security are enabled, and SPNEGO Web authentication is enabled, SPNEGO is initialized when processing a first inbound HTTP request. The Web authenticator component then interacts with SPNEGO, which is defined and enabled in the security configuration repository. When the filter criteria is met, SPNEGO is responsible for authenticating access to the secured resource that is identified in the HTTP request.

In addition to WebSphere Application Server security runtime services, some external components are required to enable the operation of SPNEGO. These external components include:

- **Windows** Microsoft Windows 2000 or Windows 2003 Servers with Active Directory domain and associated Kerberos Key Distribution Center (KDC).
- A client application, for example, Microsoft .NET, or Web service and J2EE client that supports the SPNEGO Web authentication mechanism, as defined in IETF RFC 2478. Microsoft Internet Explorer Version 5.5 or later and Mozilla Firefox Version 1.0 are browser examples. Any browser must be configured to use the SPNEGO Web authentication mechanism. For more information on performing this configuration, see “Configuring the client browser to use SPNEGO” on page 315.

The authentication of HTTP requests is triggered by the requestor (the client-side), which generates a SPNEGO token. WebSphere Application Server receives this token. Specifically, the SPNEGO Web authentication decodes and retrieves the requester’s identity from the SPNEGO token. The identity is used to establish a secure context between the requester and the application server.

SPNEGO Web authentication is a server-side solution in WebSphere Application Server. Client-side applications are responsible for generating the SPNEGO token for use by SPNEGO Web authentication. The requester’s identity in the WebSphere Application Server security registry must be identical to the identity that the SPNEGO Web authentication retrieves. An identical match does occur when Microsoft Windows Active Directory server is the Lightweight Directory Access Protocol (LDAP) server that is used in WebSphere Application Server. A custom login module is available as a plug-in to support custom mapping of the identity from the Active Directory to the WebSphere Application Server security registry.

Read about “Mapping of a client Kerberos principal name to the WebSphere user registry ID” on page 276 for more information about using this custom login module.

WebSphere Application Server validates the identity against its security registry. If the validation is successful, the client Kerberos ticket and GSS delegation credential are retrieved and placed in the client subject, which then produces a Lightweight Third Party Authentication (LTPA) security token. It then places and returns a cookie to the requester in the HTTP response. Subsequent HTTP requests from this same requester to access additional secured resources in WebSphere Application Server use the LTPA security token previously created to avoid repeated login challenges.

The Web administrator has access to the following SPNEGO security components and associated configuration data, as shown in the following figure:




Figure 10. SPNEGO Web authentication and security configuration elements

The benefits of SPNEGO Web authentication

The benefits of having WebSphere Application Server use SPNEGO as the Web authentication service for WebSphere Application Server include the following:

- **Windows** An integrated single sign-on environment with Microsoft Windows 2000 or 2003 Servers using Active Directory domain is established.
- The cost of administering a large number of ids and passwords is reduced.
- A secure and mutually authenticated transmission of security credentials from the Web browser or Microsoft .NET clients is established.
- Interoperability with Web services and Microsoft .NET, or Web service applications that use SPNEGO authentication at the transport level is achieved.
- With Kerberos authentication support, SPNEGO Web authentication can provide an end-to-end SPNEGO to Kerberos solution and preserve the Kerberos credential from the client.

SPNEGO Web authentication in a single Kerberos realm

SPNEGO Web authentication is supported in a single Kerberos realm. The challenge-response handshake process is shown in the following figure:




Figure 11. SPNEGO Web authentication in a single Kerberos realm

In the figure above, the following events occur:

1. The client sends an HTTP/Post/Get/Web-Service request to WebSphere Application Server.
2. WebSphere Application Server returns HTTP 401 Authenticate/Negotiate.
3. The client obtains a Ticket Granting Ticket (TGT).
4. The client requests a Service Ticket (TGS_REQ).
5. The client obtains a Service Ticket (TGS_REP).
6. The client sends HTTP/Post/Get/Web-Service and an authorization SPNEGO token to WebSphere Application Server.
7. WebSphere Application Server validates the SPNEGO token. If the validation is successful, it retrieves the user ID, the client Kerberos ticket and the GSS delegation credential from the SPNEGO token.
8. WebSphere Application Server validates the user ID with the WebSphere user registry and creates an LTPA token.
9. WebSphere Application Server returns HTTP 200, content and the LTPA token to the client.

Note: Other clients (for example, Web Services, .NET and J2EE) that support SPNEGO do not have to follow the challenge-response handshake process as shown above. Those clients can obtain a ticket-granting ticket (TGT) and a Kerberos service ticket for the target server, create a SPNEGO token, insert it in the HTTP header, and then follow the normal process for creating an HTTP request.

SPNEGO Web authentication in a trusted Kerberos realm

SPNEGO Web authentication is also supported in a trusted Kerberos realm. The challenge-response handshake process is shown in the following figure:

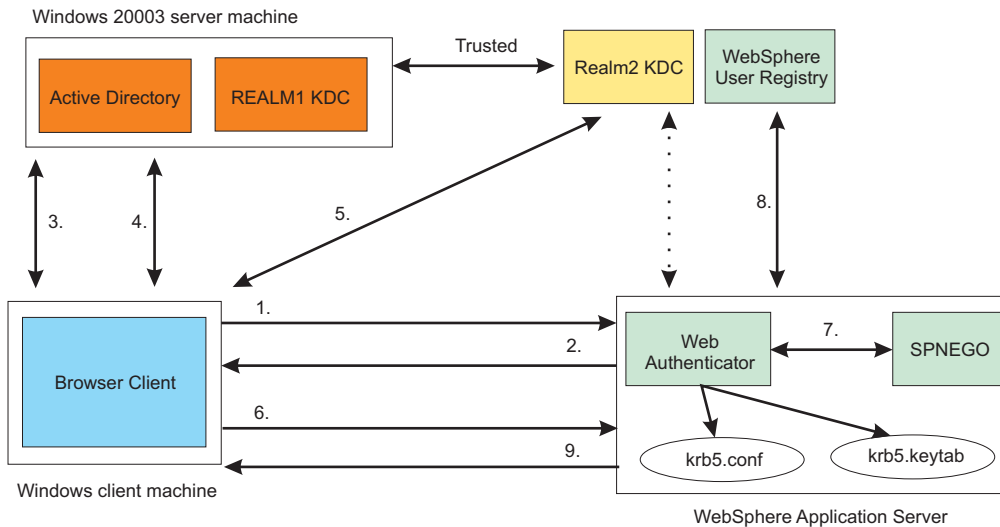


Figure 12. SPNEGO Web authentication in a trusted Kerberos realm

In the figure above, the following events occur:

1. The client sends an HTTP/Post/Get/Web-Service request to WebSphere Application Server.
2. WebSphere Application Server returns HTTP 401 Authenticate/Negotiate
3. The client obtains a Ticket Granting Ticket (TGT).
4. The client requests a cross realm ticket (TGS_REQ) for REALM2 from the REALM1 KDC.
5. The client uses the cross-realm ticket from step 4 to request a Service Ticket from the REALM2 KDC.
6. The client sends HTTP/Post/Get/Web-Service and an authorization SPNEGO token to WebSphere Application Server.
7. WebSphere Application Server validates the SPNEGO token. If the validation is successful, it retrieves the user ID, the client Kerberos ticket and the GSS delegation credential from the SPNEGO token.
8. WebSphere Application Server validates the user ID with the WebSphere user registry and creates an LTPA token.
9. WebSphere Application Server returns HTTP 200, content and the LTPA token to the client.

In the trusted Kerberos realms environment, be aware of the following:

- The Kerberos trusted realm setup must be done on each of the Kerberos KDCs. See your Kerberos Administrator and User's guide for more information about how to set up Kerberos trusted realms.
- The Kerberos client principal name from the SPNEGO token might not exist in the WebSphere user registry; the Kerberos principal mapping to the WebSphere user registry might require it.

Read about "Mapping of a client Kerberos principal name to the WebSphere user registry ID" on page 276 for more information.

Setting up SPNEGO as the Web authentication mechanism for WebSphere Application Server

Before you set up SPNEGO Web authentication in the administrative console or by using **wsadmin** commands, you must perform the following steps to set up Kerberos as the authentication mechanism for WebSphere Application Server.

Note: Kerberos authentication mechanism on the server side must be done by the system administrator. The Kerberos keytab file must be protected.

1. Ensure that the KDC is configured.
See your Kerberos Administrator and User's guide for more information.

2. The IBM Java Generic Security Service (JGSS) and KRB5 require a Kerberos configuration file (**krb5.conf** or **krb5.ini**) on each node or Java virtual machine (JVM). In this release of WebSphere Application Server, this configuration file should be placed in the `config/cells/<cell_name>` directory so that all application servers can access this file. If you do not have a Kerberos configuration file, use a **wsadmin** command to create one.

Read about “Creating a Kerberos configuration file” on page 270 for more information.

Create a Kerberos keytab file that contains the Kerberos service principal (SPN) for each SPNEGO Web authentication filter host name. The format of the SPN is `HTTP/<fully qualified hostname>`. SPNEGO uses this SPN to validate and establish the security context with the SPNEGO requester.

To use a remote HTTP server, you must create the SPN for the remote proxy server and add the proxy SPN and key to the keytab file.

The Kerberos keytab file (**krb5.keytab**) contains all of the SPNs for the node and must be protected. This file can be placed in the `config/cells/<cell_name>` directory.

Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information.

Note: SPNEGO Web authentication and Kerberos authentication both use the same Kerberos configuration and keytab files.

3. Set up SPNEGO Web authentication for WebSphere Application Server by using the administrative console.

Read about “Enabling and configuring SPNEGO Web authentication using the administrative console” for more information.

For information about how to set up SPNEGO Web authentication for WebSphere Application Server by using **wsadmin** commands, read about “SPNEGO Web authentication configuration commands” on page 1258.

Enabling and configuring SPNEGO Web authentication using the administrative console

You can enable and configure the Simple and Protected GSS-API Negotiation (SPNEGO) as the Web Authenticator for the application server by using the administrative console

Before you begin

Note: You must have completed the steps as described in “Creating a single sign-on for HTTP requests using SPNEGO Web authentication” on page 300 before enabling SPNEGO Web authentication using the administrative console.

If you do not have a Kerberos keytab file (**krb5.keytab**) then you cannot use the **createkrbConfigFile** command to create a Kerberos configuration file. Read about “Creating a Kerberos configuration file” on page 270 for more information.

You must have a Kerberos keytab file (**krb5.keytab**) that contains the Kerberos service principal name, `HTTP/<fully qualified hostname>@KerberosRealm`, for any WebSphere application server that processes an HTTP request.

1. In the administrative console, click **Security > Global security**.
2. Under Authentication, expand Web and SIP Security and then click **SPNEGO Web authentication**.

Note: You must configure the filter before enabling SPNEGO Web authentication.

3. Optional: Select the **Dynamically update SPNEGO** option if you want to dynamically update the SPNEGO run time when SPNEGO changes occur without restarting the application server.
4. Select **Enable SPNEGO** to enable the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) as a Web Authenticator for WebSphere Application Server. The **Dynamically update SPNEGO** and **Allow fall back to application authentication mechanism** options are disabled unless you select **Enable SPNEGO**.

5. Optional: When you select **Allow fall back to application authentication mechanism**, if SPNEGO authentication fails the authentication mechanism that is defined during application assembly time is used.
6. Enter the Kerberos configuration file name with its full path, or click **Browse** to locate it. The Kerberos client configuration file, **krb5.conf** or **krb5.ini**, contains Kerberos configuration information, including the locations of the Key Distribution Centers (KDCs) for the realm of interest. The **krb5.conf** file is the default name for all platforms except the Windows operating system, which uses the **krb5.ini** file.
7. Optional: Enter the Kerberos keytab file name with its full path, or click **Browse** to locate it. The Kerberos keytab file contains one or more Kerberos service principal names and keys. The default keytab file is **krb5.keytab**. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users. Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information. If you do not specify this parameter, the default Kerberos realm name in the Kerberos configuration file is used.
8. From SPNEGO filters, select an existing host name to edit or select **New** to create a new one. By convention, a Kerberos service principal name (SPN) is divided into three parts: the primary, the instance, and the Kerberos realm name. The SPNEGO service name must be HTTP, so the Kerberos service principal name for SPNEGO Web is HTTP/<fully qualified host name>@KERBEROS_REALM. The SPN is used to validate the incoming SPNEGO token and to establish security context with a requestor.
 - a. Required: On the next page, enter a fully qualified host name in the **Host name** field. The host name is part of the Kerberos service principal name (SPN), HTTP/<fully qualified host name>, used by SPNEGO to establish a Kerberos secure context. For each filter entry, the configuration code forms the Kerberos service principal as HTTP/<fully qualified host name>@KERBEROS_REALM, the Kerberos realm that you must specify in the next step. The Kerberos keytab must contain this Kerberos service principal and its keys.
 - b. Optional: In the **Kerberos realm name** field, enter the Kerberos realm name. In most cases, your realm is your domain name in uppercase letters. For example, a machine with the domain name of test.austin.ibm.com typically has a Kerberos realm name of AUSTIN.IBM.COM. If you do not specify this parameter, the default Kerberos realm name in the Kerberos configuration file is used.
 - c. Enter a filter criteria in the **Filter criteria** field. The filter criteria is the filtering parameter used by the Java class that is used by SPNEGO. It defines arbitrary criteria that is meaningful to the implementation class used.

The com.ibm.ws.security.spnego.HTTPHeaderFilter default implementation class uses this property to define a list of selection rules that represent conditions that are matched against the HTTP request headers to determine whether or not the HTTP request is selected for SPNEGO authentication.

Each condition is specified with a key-value pair, separated from each other by a semicolon. The conditions are evaluated from left to right, as they display in the specified property. If all conditions are met, the HTTP request is selected for SPNEGO authentication.

The key and value in the key-value pair are separated by an operator that defines which condition is checked. The key identifies an HTTP request header to extract from the request and its value is compared with the value that is specified in the key-value pair according to the operator specification. If the header that is identified by the key is not present in the HTTP request, the condition is treated as not being met.

Any of the standard HTTP request headers can be used as the key in the key-value pairs. Refer to the HTTP specification for the list of valid headers. In addition, two keys are defined to extract information from the request, also useful as a selection criterion, which is not available through standard HTTP request headers. The remote-address key is used as a pseudo header to retrieve the remote TCP/IP address of the client application that sent the HTTP request. The request-URL key is used as a pseudo header to retrieve the URL that is used by the client application to make the request. The interceptor uses the result of the getRequestURL operation in the

javax.servlet.http.HttpServletRequest interface to construct the Web address. If a query string is present, the result of the getQueryString operation in the same interface is also used. In this case, the complete URL is constructed as follows:

```
String url = request.getRequestURL() + '?' + request.getQueryString();
```

The following operators and conditions are defined:

Table 6. Filter conditions and operations

Condition	Operator	Example
Match exactly	= = Arguments are compared as equal.	host=host.my.company.com
Match partially (includes)	%= Arguments are compared with a partial match being valid.	user-agent%=IE 6
Match partially (includes one of many)	^= Arguments are compared with a partial match being valid for one of many arguments specified.	request-url^=webApp1 webApp2 webApp3
Does not match	!= Arguments are compared as not equal.	request-url!=noSPNEGO
Greater than	> Arguments are compared lexogaphically as greater than.	remote-address>192.168.255.130
Less than	< Arguments are compared lexogaphically as less than.	remote-address<192.168.255.135

- d. In the **Filter class** field, enter the name of the Java class that is used by SPNEGO to select which HTTP requests are subject to SPNEGO Web authentication. If you do not specify this parameter, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used.
- e. Optional: In the **SPNEGO not supported error page URL** field you can optionally enter the URL of a resource that contains the content that SPNEGO includes in the HTTP response that is displayed by the (browser) client application if it does not support SPNEGO authentication. This property can specify a Web (http://) or a file (file://) resource.

If the **SPNEGO not supported error page URL** field is not specified, or the SPNEGO that is authenticated cannot find the specified resource, the following content is used:

```
<html><head><title>SPNEGO authentication is not supported</title></head>
<body>SPNEGO authentication is not supported on this client</body></html>;
```

- f. Optional: In the **NTLM token received error page URL** field you can optionally specify the URL of a resource that contains the content that SPNEGO includes in the HTTP response, which is displayed by the browser client application. The browser client application displays this HTTP response when the browser client sends a NT LAN manager (NTLM) token instead of the expected SPNEGO token during the challenge-response handshake.

If the **NTLM token received error page URL** field is not specified, or the SPNEGO that is authenticated cannot find the specified resource, the following content is used:

```
<html><head><title>An NTLM Token was received.</title></head>
<body>Your browser configuration is correct, but you have not logged into a supported
Microsoft(R) Windows(R) Domain.
<p>Please login to the application using the normal login page.</html>
```

- g. Optional: Select **Trim Kerberos realm from principal name** to specify whether SPNEGO removes the suffix of the principal user name, starting from the @ that precedes the Kerberos realm name. If this option is selected, the suffix of the principal user name is removed. If this attribute is not selected, the suffix of the principal name is retained. The default is for this option to not be selected.
 - h. Optional: Select **Enable delegation of Kerberos credentials** to indicate whether the Kerberos delegated credentials are stored by SPNEGO. This option also enables an application to retrieve the stored credentials and to propagate them to other applications downstream for additional Kerberos authentication with the client Kerberos credential. An attempt is made to retrieve the GSS delegation credential and Kerberos tickets and to place these Kerberos credentials in the subject.
9. Click **Apply**. The filter criteria and filter class are validated if they were specified.
 10. Click **OK**. This complete the SPNEGO Web Authentication page.

Results

SPNEGO is now enabled as the Web Authenticator for the application server.

Adding or modifying SPNEGO web authentication filters using the administrative console:

The Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) filter values control different aspects of SPNEGO. You can specify different filter values for each application server by using the administrative console.

1. In the administrative console, click **Security > Global security**.
2. Under Authentication, expand Web and SIP Security and then click **SPNEGO Web authentication**.
3. Under SPNEGO filters, select an existing host name to edit or click **New**.
4. Enter a Kerberos host name if creating a new filter. This is the host name in the Kerberos service principal name (SPN) used by SPNEGO to establish a Kerberos secure context.
5. Enter a Kerberos realm name. In most cases, the realm is your domain name in uppercase letters. For example, a machine with the domain name of test.austin.ibm.com would usually have a Kerberos realm name of AUSTIN.IBM.COM
6. Enter a filter criteria in the **Filter criteria** field. The filter criteria is the filtering parameter used by the Java class that is used by SPNEGO. It defines arbitrary criteria that is meaningful to the implementation class used.

Read about “Enabling and configuring SPNEGO Web authentication using the administrative console” on page 307 for more information about the filter criteria.

7. In the **Filter class** field, enter the name of the Java class that is used by SPNEGO to select which HTTP requests are subject to SPNEGO authentication. If you do not specify this parameter, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used.
8. Optional: In the **SPNEGO not supported error page URL** field you can optionally enter the URL of a resource that contains the content which SPNEGO includes in the HTTP response that is displayed by the (browser) client application if it does not support SPNEGO authentication. This property can specify a Web (http://) or a file (file://) resource.
9. Optional: In the **NTLM token received error page URL** field you can optionally specify the URL of a resource that contains the content which SPNEGO includes in the HTTP response, which is displayed by the (browser) client application. The (browser) client application displays this HTTP response when the browser client sends a NT LAN manager (NTLM) token instead of the expected SPNEGO token during the challenge-response handshake.

10. Optional: Select **Trim Kerberos realm from principal name** to specify whether SPNEGO should remove the suffix of the principal user name, starting from the @ that precedes the Kerberos realm name. If this option is selected, the suffix of the principal user name is removed. If this attribute is not selected, the suffix of the principal name is retained. The default is for this option to be selected.
11. Optional: Select **Enable delegation of Kerberos credentials** to indicate whether the Kerberos delegated credentials should be stored by SPNEGO. This option also enables an application to retrieve the stored credentials and to propagate them to other applications downstream for additional SPNEGO authentication.
12. Click **OK**.

Results

You have modified existing or created new SPNEGO filter values for your application server.

SPNEGO Web authentication enablement:

You can enable the Simple and Protected GSS-API Negotiation (SPNEGO) as the Web Authenticator for WebSphere Application Server.

SPNEGO Web authentication provides client-server single sign-on by negotiating use of SPNEGO tokens.

To view this administrative console page, click **Security > Global security**. From Authentication, expand Web and SIP Security, and then click **SPNEGO Web authentication**.

Dynamically update SPNEGO:

Enables you to dynamically update the SPNEGO runtime when SPNEGO changes occur without restarting the application server.

Note: This option is disabled if the **Enable SPNEGO** option is not selected.

Default: Enabled

Enable SPNEGO:

Specifies the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) as a Web Authenticator for the application server.

Default: Disabled

Allow fall back to application authentication mechanism:

Specifies that SPNEGO as a Web authenticator is used to log in to WebSphere Application Server first. However, if the login fails, then the application authentication mechanism is used to log in to WebSphere Application Server.

Note: This option is disabled if the **Enable SPNEGO** option is not selected.

Default: Disabled

Kerberos configuration file with full path:

The Kerberos configuration file name with its full path. You can click **Browse** to locate it.

The Kerberos client configuration file, **krb5.conf** or **krb5.ini**, contains Kerberos configuration information, including the locations of the Key Distribution Centers (KDCs) for the realm of interest. The **krb5.conf** file is the default name for all platforms except the Windows operating system, which uses the **krb5.ini** file.

Data type: String

Kerberos keytab file name with full path:

The Kerberos keytab file name with its full path. You can click **Browse** to locate it.

The Kerberos keytab file contains one or more Kerberos service principal names and keys. The default keytab file is **krb5.keytab**. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users. Read about “Creating a Kerberos service principal and keytab file” on page 273 for more information.

If you do not specify a Kerberos keytab file then the default keytab file that is defined in the Kerberos configuration file is used.

Data type: String

SPNEGO Web authentication filter values:

The Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) Web Authentication filter values control different aspects of SPNEGO. Use this page to specify different filter values for each application server.

To view this administrative console page, click **Security > Global security**. From Authentication, expand Web and SIP Security, and then click **SPNEGO Web authentication**. Under SPNEGO Filters, click **New** or select a filter to edit.

Host name:

Specifies the fully-qualified host name in the Kerberos service principal name (SPN) that is used by SPNEGO to establish a Kerberos secure context.

The host name is the fully qualified form of the hostname. For example, myHostname.austin.ibm.com.

The Kerberos SPN is a string of the form HTTP/<fully qualified hostname>@KERBEROS_REALM . The complete SPN is used with the Java Generic Security Service (JGSS) by the SPNEGO provider to obtain the security credential and security context that are used in the authentication process.

Data type: String

Kerberos realm name:

Specifies the name of your Kerberos realm. In most cases, your realm is your domain name in uppercase letters. For example, a machine with the domain name of test.austin.ibm.com might usually have a Kerberos realm name of AUSTIN.IBM.COM.

If you do not specify the Kerberos realm name, then the default realm that is defined in the Kerberos configuration file is used.

Filter criteria:

The filtering criteria used by the Java class that is used by SPNEGO.

The `com.ibm.ws.security.spnego.HTTPHeaderFilter` default implementation class uses this property to define a list of selection rules that represent conditions that are matched against the HTTP request headers to determine whether or not the HTTP request is selected for SPNEGO authentication.

Each condition is specified with a key-value pair, separated from each other by a semicolon. The conditions are evaluated from left to right, as they display in the specified property. If all conditions are met, the HTTP request is selected for SPNEGO authentication.

The key and value in the key-value pair are separated by an operator that defines which condition is checked. The key identifies an HTTP request header to extract from the request and its value is compared with the value that is specified in the key-value pair according to the operator specification. If the header that is identified by the key is not present in the HTTP request, the condition is treated as not being met.

Any of the standard HTTP request headers can be used as the key in the key-value pairs. Refer to the HTTP specification for the list of valid headers. In addition, two keys are defined to extract information from the request, also useful as a selection criterion, which is not available through standard HTTP request headers. The `remote-address` key is used as a pseudo header to retrieve the remote TCP/IP address of the client application that sent the HTTP request. The `request-URL` key is used as a pseudo header to retrieve the URL that is used by the client application to make the request. The interceptor uses the result of the `getRequestURL` operation in the `javax.servlet.http.HttpServletRequest` interface to construct the Web address. If a query string is present, the result of the `getQueryString` operation in the same interface is also used. In this case, the complete URL is constructed as follows:

```
String url = request.getRequestURL() + '?' + request.getQueryString();
```

The following operators and conditions are defined:

Table 7. Filter conditions and operations

Condition	Operator	Example
Match exactly	= = Arguments are compared as equal.	host=host.my.company.com
Match partially (includes)	%= Arguments are compared with a partial match being valid.	user-agent%=IE 6
Match partially (includes one of many)	^= Arguments are compared with a partial match being valid for one of many arguments specified.	request-url^=webApp1 webApp2 webApp3
Does not match	!= Arguments are compared as not equal.	request-url!=noSPNEGO
Greater than	> Arguments are compared lexogaphically as greater than.	remote-address>192.168.255.130
Less than	< Arguments are compared lexogaphically as less than.	remote-address<192.168.255.135

Data type:

String

Filter class:

Specifies the name of the Java class that is used by SPNEGO to select which HTTP requests are subject to SPNEGO authentication. If you do not specify this parameter, the default filter class, `com.ibm.ws.security.spnego.HTTPHeaderFilter`, is used.

Data type: String

SPNEGO not supported error page URL:

This choice is optional. It specifies the URL of a resource that contains the content which SPNEGO includes in the HTTP response that is displayed by the browser client application if it does not support SPNEGO authentication.

This property can specify a Web (`http://`) or a file (`file://`) resource.

If this property is not specified, or the interceptor cannot find the specified resource, the following content is used:

```
<html><head><title>SPNEGO authentication is not supported</title></head>
<body>SPNEGO authentication is not supported on this client</body></html>;
```

Data type: String

NTLM token received error page URL:

This property is optional. It specifies the URL of a resource that contains the content which SPNEGO includes in the HTTP response, which is displayed by the browser client application.

The browser client application displays this HTTP response when the browser client sends an NT LAN manager (NTLM) token instead of the expected SPNEGO token during the challenge-response handshake.

If this property is not specified or the interceptor cannot find the specified resource, the following content is used:

```
<html><head><title>An NTLM Token was received.</title></head>
<body>Your browser configuration is correct, but you have not logged into a supported
Microsoft(R) Windows(R) Domain.
<p>Please login to the application using the normal login page.</html>
```

Data type: String

Enable delegation of Kerberos credentials:

Specifies whether the Kerberos delegated credentials should be stored by SPNEGO. It also enables an application to retrieve the stored credentials and to propagate them to other applications downstream for additional SPNEGO authentication.

This option requires the use of the advanced Kerberos credential delegation feature and the development of custom logic by the application developer. The developer must interact directly with the Kerberos Ticket Granting Service (TGS) to obtain a Ticket Granting Ticket (TGT) using the delegated Kerberos credentials on behalf of the user who originated the request. The developer must also construct the appropriate Kerberos SPNEGO token and include it in the HTTP request to continue the downstream SPNEGO authentication process, including handling additional SPNEGO challenge-response exchange, if necessary.

Default: Disabled

Trim Kerberos realm from principal name:

This choice is optional. It specifies whether SPNEGO removes the suffix of the principal user name, starting from the @ that precedes the Kerberos realm name. If this attribute is set to `true`, the suffix of the principal user name is removed. If this attribute is set to `false`, the suffix of the principal name is retained. The default value used is `true`.

Default: Disabled

Configuring the client browser to use SPNEGO

You can configure your browser to utilize the Simple and Protected GSS-API Negotiation (SPNEGO) mechanism.

Before you begin

You need to know how to display and set options in the Microsoft Internet Explorer browser or any other browser (such as Firefox). You must have a browser installed that supports SPNEGO authentication.

About this task

Complete the following steps to ensure that your Microsoft Internet Explorer browser is enabled to perform SPNEGO authentication.

1. At the desktop, log in to the windows active directory domain.
2. Activate Internet Explorer.
3. In the Internet Explorer window, click **Tools > Internet Options > Security** tab.
4. Select the **Local intranet** icon and click **Sites**.
5. In the Local intranet window, ensure that the "check box" to include all local (intranet) not listed in other zones is selected, then click **Advanced**.
6. In the **Local intranet** window, fill in the Add this Web site to the zone field with the Web address of the host name so that the single sign-on (SSO) can be enabled to the list Web sites shown in the Web sites field. Your site information technology staff provides this information. Click **OK** to complete this step and close the Local intranet window.
7. On the **Internet Options** window, click the **Advanced** tab and scroll to **Security settings**. Ensure that the **Enable Integrated Windows Authentication (requires restart)** box is selected.
8. Click **OK**. Restart your Microsoft Internet Explorer to activate this configuration.

Results

Complete the following steps to ensure that your Firefox browser is enabled to perform SPNEGO authentication.

1. At the desktop, log in to the windows active directory domain.
2. Activate Firefox.
3. At the address field, type **about:config**.
4. In the Filter, type **network.n**
5. Double click on **network.negotiate-auth.trusted-uris**. This preference lists the sites that are permitted to engage in SPNEGO Authentication with the browser. Enter a comma-delimited list of trusted domains or URLs.

Note: You must set the value for **network.negotiate-auth.trusted-uris**.

6. If the deployed SPNEGO solution is using the advanced Kerberos feature of Credential Delegation double click on **network.negotiate-auth.delegation-uris**. This preference lists the sites for which the browser may delegate user authorization to the server. Enter a comma-delimited list of trusted domains or URLs.
7. Click **OK**. The configuration appears as updated.
8. Restart your Firefox browser to activate this configuration.

Your Internet browser is properly configured for SPNEGO authentication. You can use applications that are deployed in WebSphere Application Server that use secured resources without being repeatedly requested for an ID and password.

Creating SPNEGO tokens for J2EE, .NET, Java, Web service clients for HTTP requests

You can create a Simple and Protected GSS-API Negotiation (SPNEGO) token for your applications and insert this token into the HTTP headers to authenticate to the WebSphere Application Server.

Before you begin

1. Create a client GSS credential. Choose one of the following 4 options:

- a. Create a GSS credential for the Kerberos credential cache. For example:

```
System.setProperty("javax.security.auth.useSubjectCredsOnly", "false");
Oid krb5MechOid = new Oid("1.2.840.113554.1.2.2");
Oid spnegoMechOid = new Oid("1.3.6.1.5.5.2");

GSSManager manager = GSSManager.getInstance();
GSSName gssUserName = manager.createName(userName, GSSName.NT_USER_NAME, krb5MechOid);
clientGssCreds = manager.createCredential(gssUserName.canonicalize(krb5MechOid),
                                         GSSCredential.INDEFINITE_LIFETIME,
                                         krb5MechOid,
                                         GSSCredential.INITIATE_ONLY);

clientGssCreds.add (gssUserName,
                   GSSCredential.INDEFINITE_LIFETIME,
                   GSSCredential.INDEFINITE_LIFETIME,
                   spnegoMechOid,
                   GSSCredential.INITIATE_ONLY);
```

- b. Create a GSS credential from a subject that has Kerberos tickets. For example:

```
Oid krb5MechOid = new Oid("1.2.840.113554.1.2.2");
Oid spnegoMechOid = new Oid("1.3.6.1.5.5.2");

GSSManager manager = GSSManager.getInstance();
clientGssCreds = (GSSCredential) Subject.doAs(subject, new PrivilegedExceptionAction()
{
    public Object run() throws GSSEException, Exception
    {
        try {
            gssName = manager.createName( userName,
                                         GSSName.NT_USER_NAME,
                                         getKrb5MechOid());
            GSSCredential gssCred = manager.createCredential(
                                     gssName.canonicalize(krb5MechOid),
                                     GSSCredential.DEFAULT_LIFETIME,
                                     krb5MechOid,
                                     GSSCredential.INITIATE_ONLY);

            gssCred.add (gssUserName,
                       GSSCredential.INDEFINITE_LIFETIME,
                       GSSCredential.INDEFINITE_LIFETIME,
                       spnegoMechOid,
                       GSSCredential.INITIATE_ONLY);

            return gssCred;
        } catch (GSSEException gsse) {
        } catch (Exception e) {
        }
    }
});
```

```

    }
    return null;
}
});

```

- c. Create a GSS credential after calling the WSKRB5Login login module. For example:

```

Oid krb5MechOid = new Oid("1.2.840.113554.1.2.2");
Oid spnegoMechOid = new Oid("1.3.6.1.5.5.2");

System.setProperty("javax.security.auth.useSubjectCredsOnly", "true");
GSSManager manager = GSSManager.getInstance();
GSSName gssUserName = manager.createName(userName, GSSName.NT_USER_NAME, krb5MechOid);
clientGssCreds = manager.createCredential(gssUserName.canonicalize(krb5MechOid),
    GSSCredential.INDEFINITE_LIFETIME,
    krb5MechOid,
    GSSCredential.INITIATE_ONLY);

clientGssCreds.add(gssUserName,
    GSSCredential.INDEFINITE_LIFETIME,
    GSSCredential.INDEFINITE_LIFETIME,
    spnegoMechOid,
    GSSCredential.INITIATE_ONLY);

```

- d. Create a GSS credential using the Microsoft native Kerberos credential cache. For example:

```

Oid krb5MechOid = new Oid("1.2.840.113554.1.2.2");
Oid spnegoMechOid = new Oid("1.3.6.1.5.5.2");

System.setProperty("javax.security.auth.useSubjectCredsOnly", "false");
GSSManager manager = GSSManager.getInstance();

clientGssCreds = manager.createCredential(null,
    GSSCredential.INDEFINITE_LIFETIME,
    krb5MechOid,
    GSSCredential.INITIATE_ONLY);

clientGssCreds.add(null,
    GSSCredential.INDEFINITE_LIFETIME,
    GSSCredential.INDEFINITE_LIFETIME,
    spnegoMechOid, GSSCredential.INITIATE_ONLY);

```

2. After you have created a client GSS credential, you can now create the SPNEGO token and insert it in the HTTP header as in the following example:

```

// create target server SPN
GSSName gssServerName = manager.createName(targetServerSpn, GSSName.NT_USER_NAME);

GSSContext clientContext = manager.createContext(gssServerName.canonicalize(spnegoMechOid),
    spnegoMechOid,
    clientGssCreds,
    GSSContext.DEFAULT_LIFETIME);

// optional enable GSS credential delegation
clientContext.requestCredDeleg(true);

byte[] spnegoToken = new byte[0];

// create a SPNEGO token for the target server
spnegoToken = clientContext.initSecContext(spnegoToken, 0, spnegoToken.length);

URL url = new URL(targetUrl);
URLConnection con= (URLConnection) url.openConnection();

try {
    // insert SPNEGO token in the HTTP header
    con.setRequestProperty("Authorization", "Negotiate " + Base64.encode(spnegoToken));
    con.getResponseCode();
} catch (IOException e) {
} catch (Exception ex) {
}
}

```

Results

Your application might need a Kerberos configuration file (**krb5.ini** or **krb5.conf**). Read about “Creating a Kerberos configuration file” on page 270 for more information.

Creating a single sign-on for HTTP requests using the SPNEGO TAI (deprecated)

Creating single sign-ons for HTTP requests using the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server requires the performance of several distinct, yet related functions that when completed, allow HTTP users to log in and authenticate only once at their desktop and receive automatic authentication from the WebSphere Application Server.

Before you begin

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Before starting this task, complete the following checklist:

- **Windows** A Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC).
- **Windows** A Microsoft Windows 2000 or Windows 2003 domain member (client) for example, a browser or Microsoft .NET client, that supports the SPNEGO authentication mechanism, as defined in IETF RFC 2478. Microsoft Internet Explorer Version 5.5 or later and Mozilla Firefox Version 1.0 qualify as such clients.

Note: A running domain controller and at least one client machine in that domain is required. Trying to use SPNEGO directly from the domain controller is not supported

- The domain member has users who can log on to the domain. Specifically, you need to have a functioning Microsoft Windows 2000 or Windows 2003 active directory domain that includes:
 - Domain controller
 - Client workstation
 - Users who can login to the client workstation
- A server platform with WebSphere Application Server running and application security enabled.
- Users on the active directory must be able to access WebSphere Application Server protected resources using a native WebSphere Application Server authentication mechanism.
- The domain controller and the host of WebSphere Application Server should have the same local time.
- Ensure the clock on clients, Microsoft Active Directory and WebSphere Application Server are synchronized to within five minutes.
- Be aware that client browsers have to be SPNEGO enabled, which you perform on the client application machine (with details explained in step 2 of this task).

About this task

The objective of this machine arrangement is to permit users to successfully access WebSphere Application Server resources without having to reauthenticate and thus achieve Microsoft Windows desktop single sign-on capability.

Configuring the members of this environment to establish Microsoft Windows single sign-on involves specific activities that are performed on three distinct machines:

- Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC)
- A Microsoft Windows 2000 or Windows 2003 domain member (client application), such as a browser or Microsoft .NET client.
- A server platform with WebSphere Application Server running.

Perform the following steps on the indicated machines to create single sign-on for HTTP requests using SPNEGO

1. **Domain Controller Machine** - Configure the Microsoft Windows 2000 or Windows 2003 Server running the Active Directory Domain Controller and associated Kerberos Key Distribution Center (KDC) This configuration activity has the following steps:

- Create a user account for the WebSphere Application Server in a Microsoft Active Directory. This account will be eventually mapped to the Kerberos service principal name (SPN).
- On the Microsoft Active Directory machine where the Kerberos key distribution center (KDC) is active, map the user account to the Kerberos service principal name (SPN). This user account represents the WebSphere Application Server as being a Kerberize'd service with the KDC. Use the **setspn** command to map the Kerberos service principal name to a Microsoft user account. The topic, "Creating a Kerberos service principal and keytab file that is used by the WebSphere Application Server SPNEGO TAI (deprecated)" on page 324 has more details about using the **setspn** command.
- Create the Kerberos keytab file and make it available to WebSphere Application Server. Use the **ktpass** tool to create the Kerberos keytab file (krb5.keytab). The topic, "Creating a Kerberos service principal and keytab file that is used by the WebSphere Application Server SPNEGO TAI (deprecated)" on page 324 has more details about using the **ktpass** command. to create the keytab file.

Note: You make the keytab file available to WebSphere Application Server by copying the krb5.keytab file from the Domain Controller (LDAP machine) to the WebSphere Application Server machine. See "Using the ktab command to manage the Kerberos keytab file" on page 327 for more details.

Note: Your domain controller operations must lead to the following results:

- A user account is created in the Microsoft Active Directory and mapped to a Kerberos service principal name.
 - A Kerberos keytab file (krb5.keytab) is created and made available to the WebSphere Application Server. The Kerberos keytab file contains the Kerberos service principal keys WebSphere Application Server uses to authenticate the user in the Microsoft Active Directory and the Kerberos account.
2. **Client Application Machine** - Configure the client application. Client-side applications are responsible for generating the SPNEGO token for use by the SPNEGO TAI. You begin this configuration process by configuring your Web browser to use SPNEGO authentication. See "Configuring the client browser to use SPNEGO TAI (deprecated)" on page 341 for the detailed steps required for your browser.
 3. **WebSphere Application Server Machine** - Configure and enable the Application Server and the associated SPNEGO TAI by performing the following tasks:
 - Ensure that LTPA is enabled. See "Configuring the Lightweight Third Party Authentication mechanism" on page 250 for more details.
 - Enable the SPNEGO TAI. See "Configuring WebSphere Application Server and enabling the SPNEGO TAI (deprecated)" on page 328 for more details.
 - Create SPNEGO TAI properties using either the wsadmin command task or the administrative console.

- For using the wsadmin command task, see
 - “SpnegoTAICommands group for the AdminTask object (deprecated)” on page 1250
 - For using the administrative console, see “Configuring WebSphere Application Server and enabling the SPNEGO TAI (deprecated)” on page 328 for more details.
 - Configure JVM properties and enable the SPNEGO TAI in Application Server in which it is defined. See “Configuring JVM custom properties, filtering HTTP requests, and enabling SPNEGO TAI in WebSphere Application Server (deprecated)” on page 342 or “Enabling the SPNEGO TAI as JVM custom property using scripting (deprecated)” on page 343 for more details.
 - Install the Kerberos keytab file (created in step 1) on the WebSphere Application Server machine. “Creating a Kerberos service principal and keytab file that is used by the WebSphere Application Server SPNEGO TAI (deprecated)” on page 324 provides the details.
 - Create a basic Kerberos configuration file (krb5.ini or krb5.conf). See “The Kerberos configuration file” on page 1255 for details.
 - Map the client Kerberos principal name to the WebSphere user registry ID, but only if the WebSphere Application Server does not use Microsoft Active Directory. See “Mapping Kerberos client principal name to WebSphere user registry ID for SPNEGO TAI (deprecated)” on page 347 for more details.
4. Optional: **Using a remote HTTP server** - To use a remote server, you must complete the following steps, which assume that you have already configured the JVM properties and enabled the SPNEGO TAI in the Application Server in which it is defined (as described in the previous three steps).
- a. Complete the steps in “Creating a Kerberos service principal and keytab file that is used by the WebSphere Application Server SPNEGO TAI (deprecated)” on page 324 for the remote proxy server.
 - b. Merge the previous keytab file created in step 1 with the keytab file created in step 4a. See “Using the ktab command to manage the Kerberos keytab file” on page 327 for more information.
 - c. Create the SPN for the remote proxy server using the addSpnegoTAIProperties wsadmin command task. For more information, see “SpnegoTAICommands group for the AdminTask object (deprecated)” on page 1250.
 - d. Restart the WebSphere Application Server.

Single sign-on for HTTP requests using SPNEGO TAI (deprecated)

WebSphere Application Server provides a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources in WebSphere Application Server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Read about “Creating a single sign-on for HTTP requests using SPNEGO Web authentication” on page 300 for more information.

SPNEGO is a standard specification defined in The Simple and Protected GSS-API Negotiation Mechanism (IETF RFC 2478).

When WebSphere Application Server administrative security is enabled, the SPNEGO TAI is initialized. While processing inbound HTTP requests, the Web authenticator component interacts with the SPNEGO

TAI, which is defined and enabled in the security configuration repository. One interceptor is selected and is responsible for authenticating access to the secured resource that is identified in the HTTP request.

Note: The use of TAIs is an optional feature. If no TAI is selected, the authentication process continues normally.

HTTP users log in and authenticate only once at their desktop and are subsequently authenticated (internally) with WebSphere Application Server. The SPNEGO TAI is invisible to the end-user of WebSphere applications. The SPNEGO TAI is only visible to the Web administrator who is responsible for ensuring a proper configuration, capacity, and maintenance of the Web environment.

In addition to WebSphere Application Server security runtime services, some external components are required to completely enable operation of the SPNEGO TAI. The external components include:

- **Windows** Microsoft Windows 2000 or Windows 2003 Servers with Active Directory domain and associated Kerberos Key Distribution Center (KDC).
- A client application, for example, a browser or Microsoft .NET client, that supports the SPNEGO authentication mechanism, as defined in IETF RFC 2478. Microsoft Internet Explorer Version 5.5 or later and Mozilla Firefox Version 1.0 are browser examples. Any browser needs to be configured to use the SPNEGO mechanism. For more information on performing this configuration, see “Configuring the client browser to use SPNEGO TAI (deprecated)” on page 341.

The authentication of HTTP requests is triggered by the requestor (the client-side), which generates a SPNEGO token. WebSphere Application Server receives this token and validates trust between the requester and WebSphere Application Server. Specifically, the SPNEGO TAI decodes and retrieves the requester’s identity from the SPNEGO token. The identity is used to establish a secure context between the requester and the application server.

Note: The SPNEGO TAI is a server-side solution in WebSphere Application Server. Client-side applications are responsible for generating the SPNEGO token for use by the SPNEGO TAI. The requester’s identity in WebSphere Application Server security registry must be identical to that identity the SPNEGO TAI retrieves. An identical match does occur when Microsoft Windows Active Directory server is the Lightweight Directory Access Protocol (LDAP) server that is used in WebSphere Application Server. A custom login module is available as a plug-in to support custom mapping of the identity from the Active Directory to the WebSphere Application Server security registry. See “Mapping Kerberos client principal name to WebSphere user registry ID for SPNEGO TAI (deprecated)” on page 347 for details on using this custom login module.

WebSphere Application Server validates the identity against its security registry and, if the validation is successful, produces a Lightweight Third Party Authentication (LTPA) security token and places and returns a cookie to the requester in the HTTP response. Subsequent HTTP requests from this same requester to access additional secured resources in WebSphere Application Server use the LTPA security token previously created, to avoid repeated login challenges.

The challenge-response handshake process is illustrated in the following graphic:

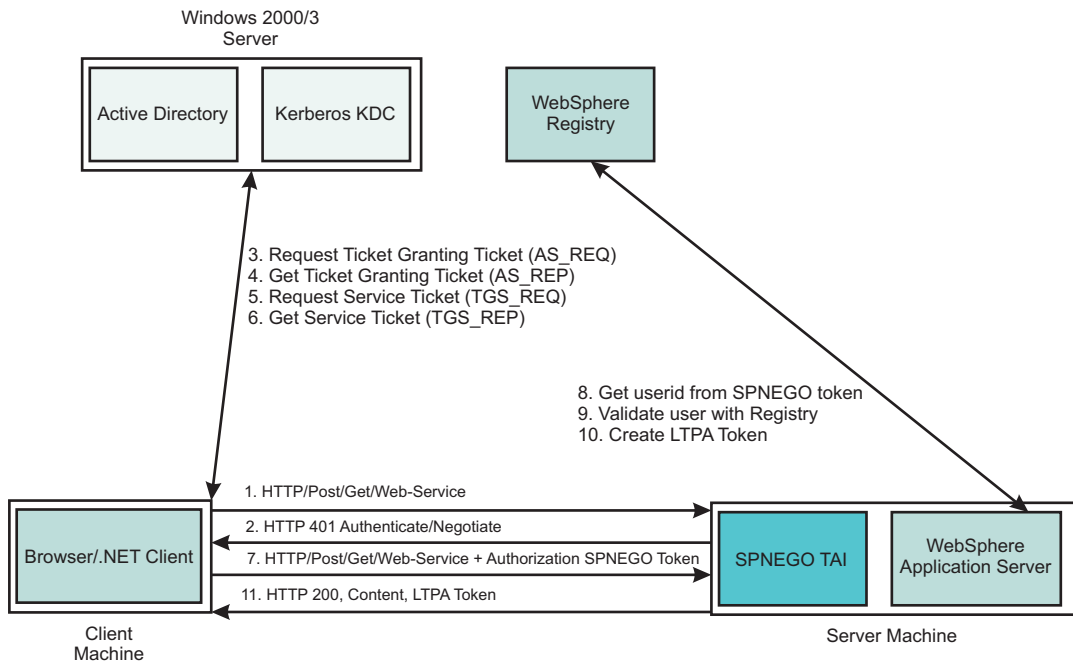


Figure 13. HTTP request processing, WebSphere Application Server - SPNEGO TAI

The SPNEGO TAI can be enabled for all or for selected WebSphere Application Servers in a WebSphere Application Server cell configuration. Also, the behavior of each SPNEGO TAI instance is controlled by custom configuration properties that are used to identify, for example, the criteria used to filter HTTP requests, such as the host name and security realm name used to construct the Kerberos Service Principal Name (SPN). For more information regarding establishing and setting the SPNEGO TAI custom configuration properties, see the following topics:

- Setting up the Kerberos configuration properties. See “The Kerberos configuration file” on page 1255.
- Setting or adjusting the SPNEGO TAI custom properties. See “SPNEGO TAI custom properties configuration (deprecated)” on page 336.
- Adjusting the SPNEGO TAI filter settings. See “Configuring JVM custom properties, filtering HTTP requests, and enabling SPNEGO TAI in WebSphere Application Server (deprecated)” on page 342
- Using the custom login module to map the identity from the Active Directory to the WebSphere Application Server registry. See Mapping user Ids from client to server for SPNEGO.
- Setting the major and additional Java virtual machine (JVM) custom properties. See “SPNEGO TAI JVM configuration custom properties (deprecated)” on page 344

The Web administrator has access to the following SPNEGO TAI security components and associated configuration data, as illustrated in the following graphic.

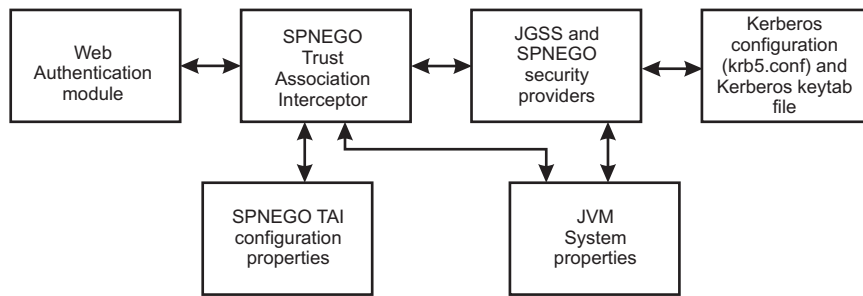


Figure 14. SPNEGO TAI security and configuration elements

- The Web authentication module and the Lightweight Third Party Authentication (LTPA) mechanism provide the plug-in runtime framework for trust association interceptors. See “Configuring the Lightweight Third Party Authentication mechanism” on page 250 for more detail is configuring the LTPA mechanism for use with the SPNEGO TAI.
- The Java Generic Security Service (JGSS) provider is included in the Java SDK (`jre/lib/ibmjgssprovider.jar`) and used to obtain the Kerberos security context and credentials that are used for authentication. IBM JGSS 1.0 is a Java Generic Security Service Application Programming Interface (GSSAPI) framework with Kerberos V5 as the underlying default security mechanism. GSSAPI is a standardized abstract interface under which can be plugged different security mechanisms based on private-key, public-key and other security technologies. GSSAPI shields secure applications from the complexities and peculiarities of the different underlying security mechanisms. GSSAPI provides identity and message origin authentication, message integrity, and message confidentiality. For more information, see JGSS.
- The Kerberos configuration properties (`krb5.conf` or `krb5.ini`) and Kerberos encryption keys (stored in a Kerberos keytab file) are used to establish secure mutual authentication.
The Kerberos key table manager (Ktab), which is part of JGSS, allows you to manage the principal names and service keys stored in a local Kerberos keytab file. Principal name and key pairs listed in the Kerberos keytab file allow services running on a host to authenticate themselves to the Kerberos Key Distribution Center (KDC). Before a server can use Kerberos, a Kerberos keytab file must be initialized on the host that runs the server.
“Using the `ktab` command to manage the Kerberos keytab file” on page 327 highlights the Kerberos configuration requirements for the SPNEGO TAI as well as the use of Ktab.
- The SPNEGO provider supplies the implementation of the SPNEGO authentication mechanism, located at `/$WAS_HOME/java/jre/lib/ext/ibmspnego.jar`.
- The custom configuration properties control the runtime behavior of the SPNEGO TAI. Configuration operations are performed with the administrative console or scripting facilities. Refer to “SPNEGO TAI custom properties configuration (deprecated)” on page 336 for more information about these custom configuration properties.
- Java virtual machine (JVM) custom properties control diagnostic trace information for problem determination of the JGSS security provider and use of the property reload feature. “SPNEGO TAI JVM configuration custom properties (deprecated)” on page 344 describes these JVM custom properties

The benefits of having WebSphere Application Server use the SPNEGO TAI include:

- **Windows** An integrated single signon environment with Microsoft Windows 2000 or 2003 Servers using Active Directory domain is established.
- The cost of administering a large number of ids and passwords is reduced.
- A secure and mutually authenticated transmission of security credentials from the Web browser or Microsoft .NET clients is established.
- Interoperability with Web services and Microsoft .NET applications that use SPNEGO authentication at the transport level is achieved.

Using the SPNEGO TAI in your WebSphere Application Server environment requires planning then implementation. See “Single sign-on capability with SPNEGO TAI - checklist (deprecated)” on page 351 in planning for SPNEGO TAI. Implementing the use of the SPNEGO TAI is divided into the following areas of responsibility:

End browser user

The end user must configure the Web browser or Microsoft .NET application to issue HTTP requests that are processed by the SPNEGO TAI.

Web administrator

The Web administrator is responsible for configuring the SPNEGO TAI of WebSphere Application Server to respond to HTTP requests of the client.

WebSphere Application Server administrator

The WebSphere Application Server administrator is responsible for configuring WebSphere Application Server and the SPNEGO TAI for optimum installation performance.

See “Creating a single sign-on for HTTP requests using the SPNEGO TAI (deprecated)” on page 318 for an explanation of the tasks required to use the SPNEGO TAI and how the responsible party performs these tasks.

Creating a Kerberos service principal and keytab file that is used by the WebSphere Application Server SPNEGO TAI (deprecated)

You perform this configuration task on the Microsoft Active Directory domain controller machine. This task is a necessary part of preparing to process single sign on browser requests to WebSphere Application Server and the SPNEGO trust association interceptor (TAI).

Before you begin

You need to have a running domain controller and at least one client machine in that domain.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

This task is performed on the active directory domain controller machine. Complete the following steps to ensure that the Microsoft Windows 2000 or Windows 2003 Server that is running the active directory domain controller is configured properly to the associated key distribution center (KDC).

1. Create a user account in the Microsoft Active Directory for the WebSphere Application Server.

Click **Start->Programs->Administrative Tools->Active Directory Users and Computers**

Use the name for the WebSphere Application Server. For example, if the Application Server you are running on the WebSphere Application Server machine is called `myappserver.austin.ibm.com`, create a new user in Active Directory called `myappserver`.

Note: Do not select “User must change password at next logon.”

Note: Make sure that you do not have the computer name `myappserver` under Computers and Domain Controllers (You check for this condition as illustrated below.). If you already have a computer name `myappserver`, then you need to create a different user account name.

- Goto **Start -> Programs -> Administrative Tools -> Active Directory Users and Computers->Computers**
- Goto **Start -> Programs -> Administrative Tools -> Active Directory Users and Computers->Domain Controllers**

2. Use the **setspn** command to map the Kerberos service principal name, HTTP/<host name>, to a Microsoft user account. An example of **setspn** usage is as follows:

```
C:\Program Files\Support Tools>
setspn -A HTTP/myappserver.austin.ibm.com myappserver
```

Note: There may already be some SPNs related to the Microsoft Windows hosts that have been added to the domain. You can display those that exist by using the **setspn -L** command, but you still have to add an HTTP SPN for WebSphere Application Server. For example, **setspn -L myappserver** would list the SPNs.

Note: Make sure that you do not have the same SPNs mapping to more than one Microsoft user account. If you map the same SPN to more than one user account, the web browser client can send a NTLM instead of SPNEGO token to WebSphere Application Server.

More information about the **setspn** command can be found here, Windows 2003 Technical Reference (setspn command)

3. Create the Kerberos keytab file and make it available to WebSphere Application Server. Use the **ktpass** command to create the Kerberos keytab file (krb5.keytab).

Use the **ktpass** tool from the Windows Server toolkit to create the Kerberos keytab file for the service principal name (SPN). Use the **ktpass** tool that matches the Windows server level you are using. That is, use the Windows 2000 version for a Windows 2000 Server, or a Windows 2003 version for a Windows 2003 server. The Windows 2003 server version of the **ktpass** tool supports the encryption type, RC4-HMAC, and Single data encryption standard (DES). The Windows 2000 server version of the **ktpass** tool are similar, but different options are necessary for the RC4-HMAC encryption type and single DES. For more information about the **ktpass** tool, see Windows 2003 Technical Reference (Kerberos keytab file and ktpass command).

Below is a summary of the functions available when you enter **ktpass -?** on the command line.

```
C:\MS SDK>ktpass -?
```

Command line options:

```
-----most useful args
[- /]      out : Keytab to produce
[- /]      princ : Principal name (user@REALM)
[- /]      pass : password to use
           use "*" to prompt for password.
-----less useful stuff
[- /]      mapuser : map princ (above) to this user account (default: don't)
[- /]      mapOp : how to set the mapping attribute (default: add it)
[- /]      mapOp : is one of:
[- /]      mapOp :      add : add value (default)
[- /]      mapOp :      set : set value
[- +]      DesOnly : Set account for des-only encryption (default:no)
[- /]      in : Keytab to read/digest
-----options for key generation
[- /]      crypto : Cryptosystem to use
[- /]      crypto : is one of:
[- /]      crypto : DES-CBC-CRC : for compatibility
[- /]      crypto : DES-CBC-MD5 : default
[- /]      crypto :      RC4 :
[- /]      ptype : principal type in question
[- /]      ptype : is one of:
[- /]      ptype : KRB5_NT_PRINCIPAL : The general ptype-- recommended
[- /]      ptype : KRB5_NT_SRV_INST : user service instance
[- /]      ptype : KRB5_NT_SRV_HST : host service instance
[- /]      kvno : Override Key Version Number
```


Default: query DC for kvno. Use /kvno 1 for Win2K compat.
 [- +] Answer : +Answer answers YES to prompts. -Answer answers NO.
 [- /] Target : Which DC to use. Default:detect

Depending on the encryption type, you use the **ktpass** tool in one of the following ways to create the Kerberos keytab file:

Note: Do not use the **-pass** switch on the **ktpass** command to reset a password for a Microsoft Windows server account.

See Windows 2003 Technical Reference (Kerberos keytab file and ktpass command) for more information.

- For a single DES encryption type

From a command prompt, run the **ktpass** command:

```
ktpass -out c:\temp\myappserver.keytab
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM
-mapUser myappserv
-mapOp set -pass was1edu
-crypto DES-CBC-MD5
+DesOnly
```

Table 8. Using ktpass for a single DES encryption type

Option	Explanation
-out c:\temp\myappserver.keytab	The key is written to this output file.
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM	The concatenation of the user logon name, and the realm must be in uppercase.
-mapUser	The key is mapped to the user, myappserver.
-mapOp	This option sets the mapping.
-pass was1edu	This option is the password for the user ID.
-crypto DES-CBC-MD5	This option uses the single DES encryption type.
+DesOnly	This option generates only DES encryptions.

- For the RC4-HMAC encryption type

Note: RC4-HMAC encryption is only supported when using a Windows 2003 Server as KDC. RC4-HMAC encryption is not supported with a Windows 2000 Server as KDC.

From a command prompt, run the **ktpass** command.

```
ktpass -out c:\temp\myappserver.keytab
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM
-mapUser myappserver
-mapOp set
-pass was1edu
-crypto RC4-HMAC
```

Table 9. Using ktpass for the RC4-HMAC encryption type

Option	Explanation
-out c:\temp\myappserver.keytab	The key is written to this output file.
-princ HTTP/myappserver.austin.ibm.com@WSSEC.AUSTIN.IBM.COM	The concatenation of the user logon name, and the realm must be in uppercase.
-mapUser	The key is mapped to the user, myappserver.
-mapOp	This option sets the mapping.
-pass was1edu	This option is the password for the user ID.
-crypto RC4-HMAC	This option chooses the RC4-HMAC encryption type.

The Kerberos keytab file is created for use with the SPNEGO TAI.

Note: A Kerberos keytab configuration file contains a list of keys that are analogous to user passwords. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk, which makes them readable only by authorized users.

You make the keytab file available to WebSphere Application Server by copying the krb5.keytab file from the Domain Controller (LDAP machine) to the WebSphere Application Server machine.

```
ftp> bin
ftp> put c:\temp\KRB5_NT_SEV_HST\krb5.keytab
```

Results

Your active directory domain controller is properly configured to process single sign on requests to WebSphere Application Server and the SPNEGO TAI

Using the ktab command to manage the Kerberos keytab file:

The Kerberos key table manager command (Ktab) allows the WebSphere administrator to manage the Kerberos service principal names and keys stored in a local Kerberos keytab file. In JDK 1.6 or later, you can use the ktab command to merge two Kerberos keytab files.

Kerberos service principal (SPN) name and keys listed in the Kerberos keytab file allow services running on the host to validate the incoming Kerberos or SPNEGO token request. Before configuring Kerberos or SPNEGO web authentication, the WebSphere Application Server administrator must setup a Kerberos keytab file on the host that is running WebSphere Application Server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated.

SPNEGO Web authentication has taken its place to provide the following enhancements:

- Configure and enable SPNEGO Web Authentication and filters on the WebSphere Application server side by using the administrative console.
- Provide dynamic reload of SPNEGO without having to stop and restart the WebSphere Application server.
- Provide fallback to an application login method if the SPNEGO Web authentication fails.

Note:

- It is important to protect the keytab files, making them readable only by authorized WebSphere users.
- Any updates to the Kerberos keytab file using Ktab do not affect the Kerberos database. If you change the keys in the Kerberos keytab file, you must also make the corresponding changes to the Kerberos database.

The syntax of Ktab is illustrated below by using Ktab with the `-help` operand.

```
$ ktab -help
```

```
Usage: java com.ibm.security.krb5.internal.tools.Ktab [options]
Available options:
-l                               list the keytab name and entries
-a <principal_name> [password]  add an entry to the keytab
```

-d <principal_name> delete an entry from the keytab
-k <keytab_name> specify keytab name and path with FILE: prefix
-m <keytab_name> <keytab_name> specify merging source keytab file name and destination keytab file name

Below is an example of how Ktab is used to merge **krb5Host1.conf** to the **krb5.conf** file:

```
[root@wssecjibe bin]# ./ktab -m /etc/krb5Host1.conf /etc/krb5.conf
Merging keytab files: source=krb5Host1.conf destination=krb5.conf
Done!
[root@wssecjibe bin]# ls /etc/krb5.*
/etc/krb5Host1.conf/etc/krb5.conf
/etc/krb5.keytab
```

Below is an example of how Ktab is used on a LINUX platform to add new principal names to the Kerberos keytab file, where ot56prod is the password for the Kerberos principal name:

```
[root@wssecjibe bin]# ./ktab -a
HTTP/wssecjibe.austin.ibm.com@WSSEC.AUSTIN.IBM.COM ot56prod -k /etc/krb5.keytab
Done!
Service key for principal HTTP/wssecjibe.austin.ibm.com@WSSEC.AUSTIN.IBM.COM saved
```

Below is an example of how Ktab is used on a Linux platform to list Kerberos keytab file content.

```
[root@wssecjibe bin]# ./ktab

KVNO Principal
----
1 HTTP/wssecjibe.austin.ibm.com@WSSEC.AUSTIN.IBM.COM
```

```
[root@wssecjibe bin]# ls /etc/krb5.*
/etc/krb5.conf
/etc/krb5.keytab
```

Note: On WebSphere Application Server, Ktab is located at:

```
<WAS install root>/java/jre/bin
```

Configuring WebSphere Application Server and enabling the SPNEGO TAI (deprecated)

Performing this task helps you, as Web administrator, to ensure that WebSphere Application Server is properly configured to enable the operation of the Simple and Protected GSS-API Negotiation (SPNEGO) trust association interceptor (TAI).

Before you begin

You need to know how to use the WebSphere Application Server administrative console to manage the security configuration and have the proper authority to modify the security configuration of the application server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

Complete the following steps to enable the operation of the SPNEGO TAI.

1. Log on to the WebSphere Application Server administrative console.
2. Click **Security > Global security**.
3. Expand **Web security** and click **Trust association**.
4. Under the General Properties heading, select the **Enable trust association** check box, then click **Interceptors**.
5. Select the SPNEGO TAI in the list of interceptors.
6. Then click **Custom properties**.
7. Click **New** and then fill in the **Name** and **Value** fields. Click **OK**. Repeat this step for each custom property that you want to apply to the SPNEGO TAI. See “SPNEGO TAI custom properties configuration (deprecated)” on page 336 for a complete list of SPNEGO TAI custom properties.

Note: It is recommended that you use the **wsadmin** utility to manage the SPNEGO TAI properties. You can add, modify, and delete SPNEGO TAI properties as well as display them using **wsadmin**. See “Adding SPNEGO TAI properties using the **wsadmin** utility (deprecated)” to add, “Modifying SPNEGO TAI properties using the **wsadmin** utility (deprecated)” on page 331 to modify, and “Deleting SPNEGO TAI properties using the **wsadmin** utility (deprecated)” on page 334 to delete SPNEGO TAI properties.

8. After you finish defining your custom properties, click **Save** to store the updated SPNEGO TAI configuration.

Results

Your SPNEGO TAI configuration is now configured for WebSphere Application Server.

Adding SPNEGO TAI properties using the wsadmin utility (deprecated):

You use the **wsadmin** utility to add properties for the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) in the security configuration for WebSphere Application Server.

About this task

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Verify that end-user desktop browsers are configured to support SPNEGO authentication, that the SPNEGO TAI is enabled, that the Java virtual machine (JVM) property is set and, that WebSphere Application Server is configured to enable the operation of the SPNEGO TAI.

Use the **wsadmin** utility to configure the SPNEGO TAI for WebSphere Application Server:

1. Start WebSphere Application Server.
2. Start the command-line utility by running the **wsadmin** command from the `app_server_root/bin` directory.
3. At the **wsadmin** prompt, enter the following command:

\$AdminTask addSpnegoTAIProperties

You can use the following parameters with this command:

Option	Description
<spnId>	This parameter is optional. It is the SPN identifier for the group of custom properties that are to be defined with this command. If you do not specify this parameter, an unused SPN identifier is assigned.
<host>	This parameter is required. It specifies the host name portion in the SPN used by the SPNEGO TAI to establish a Kerberos secure context.
<filter>	This parameter is optional. It defines the filtering criteria used by the class specified with the above attribute. If you do not specify this parameter, all HTTP requests are subject to SPNEGO authentication.
<filterClass>	This parameter is optional. It specifies the name of the Java class used by the SPNEGO TAI to select which HTTP requests will be subject to SPNEGO authentication. If you do not specify this parameter, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used.
<noSpnegoPage>	<p>This parameter is optional. It specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application if it does not support SPNEGO authentication.</p> <p>If you do not specify the noSpnegoPage parameter then the default is used:</p> <pre>"<html><head><title>SPNEGO authentication is not supported.</title></head>" + "<body>SPNEGO authentication is not supported on this client.</body></html>";</pre>
<ntlmTokenPage>	<p>This parameter is optional. It specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response that is to be displayed by the (browser) client application when the SPNEGO token received by the interceptor (after the challenge-response handshake) contains a NT LAN manager (NTLM) token instead of the expected SPNEGO token.</p> <p>If you do not specify the ntlmTokenPage parameter then the default is used:</p> <pre>"<html><head><title>An NTLM Token was received.</title></head>" + "<body>Your browser configuration is correct, but you have not logged into a supported Windows Domain." + "<p>Please login to the application using the normal login page.</html>";</pre>

Option	Description
<trimUserName>	This parameter is optional. It specifies whether or not the SPNEGO TAI is to remove the suffix of the principal user name, starting from the "@" that precedes the Kerberos realm name. If this parameter is set to true, the suffix of the principal user name is removed. If this parameter is set to false, the suffix of the principal name is retained. The default value used is true.

Results

SPNEGO TAI properties have been added for this WebSphere Application Server.

Example

Example 1

The following example configures the SPNEGO TAI to intercept HTTP requests that contain IE 6 in the user agent request header. The SPNEGO TAI uses the SPN of HTTP/myhost.ibm.com@<default_realm> to authenticate the request originator.

```
$AdminTask addSpnegoTAIProperties -host myhost.ibm.com -filter user-agent%=IE 6
```

Example 2

The following is an example of adding SPNEGOTAIProperties for SPN1 to use the default filterClass and to intercept all requests for the host, central01.austin.ibm.com.

```
wsadmin>$AdminTask addSpnegoTAIProperties -interactive
Add SPNEGO TAI properties
```

Add SPNEGO TAI configuration properties.

```
*Host name in Service Principal Name (host): central01.austin.ibm.com
Service Principal Name identifier (spnId): 1
HTTP header filter rule (filter):
Name of class used to filter HTTP requests (filterClass):
SPNEGO not supported browser response (noSpnegoPage):
NTLM Token received browser response (ntlmTokenPage):
Trim User Name browser response (trimUserName):
```

Add SPNEGO TAI properties

```
F (Finish)
C (Cancel)
```

```
Select [F, C]: [F] f
```

```
WASX7278I: Generated command line: $AdminTask addSpnegoTAIProperties {-host central01.austin.ibm.com}
com.ibm.ws.security.spnego.SPN1.hostName=central01.austin.ibm.com
wsadmin>
```

Modifying SPNEGO TAI properties using the wsadmin utility (deprecated):

You use the wsadmin utility to modify the properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

About this task

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and

authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Verify that end-user desktop browsers are configured to support SPNEGO authentication, that the SPNEGO TAI is enabled, that the Java virtual machine (JVM) property is set and, that WebSphere Application Server is configured to enable the operation of the SPNEGO TAI.

You use the `wsadmin` utility to configure the SPNEGO TAI for WebSphere Application Server:

1. Start WebSphere Application Server.
2. Start the command-line utility by running the `wsadmin` command from the `app_server_root/bin` directory.
3. At the `wsadmin` prompt, enter the following command:
`$AdminTask modifySpnegoTAIProperties`

You can use the following parameters with this command:

Option	Description
<code><spnId></code>	This parameter is required. It is the SPN identifier for the group of custom properties that are to be defined with this command.
<code><host></code>	This parameter is optional. It specifies the host name portion in the SPN used by the SPNEGO TAI to establish a Kerberos secure context.
<code><filter></code>	This parameter is optional. It defines the filtering criteria used by the class specified with the above attribute.
<code><filterClass></code>	This parameter is optional. It specifies the name of the Java class used by the SPNEGO TAI to select which HTTP requests will be subject to SPNEGO authentication. If no class is specified, all HTTP requests will be subject to SPNEGO authentication.
<code><noSpnegoPage></code>	This parameter is optional. It specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application if it does not support SPNEGO authentication. If you do not specify the <code>noSpnegoPage</code> attribute then the default is used: <pre>"<html><head><title>SPNEGO authentication is not supported. </title></head>" + "<body>SPNEGO authentication is not supported on this client. </body></html>";</pre>

Option	Description
<ntlmTokenPage>	<p>This parameter is optional. The ntlmTokenPage parameter specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response, which will be displayed by the (browser) client application. The (browser) client application displays this HTTP response when the browser client sends a NT LAN manager (NTLM) token instead of the expected SPNEGO token during the challenge-response handshake.</p> <p>If you do not specify the ntlmTokenPage attribute then the default is used:</p> <pre>"<html><head><title>An NTLM Token was received.</title></head>" + "<body>Your browser configuration is correct, but you have not logged into a supported Windows Domain." + "<p>Please login to the application using the normal login page.</html>";</pre>
<trimUserName>	<p>This parameter is optional. It specifies whether (true) or not (false) the SPNEGO TAI is to remove the suffix of the principal user name, starting from the "@" that precedes the Kerberos realm name. If this attribute is set to true, the suffix of the principal user name is removed. If this attribute is set to false, the suffix of the principal name is retained. The default value used is true.</p>

Results

SPNEGO TAI properties are modified for this WebSphere Application Server.

Example

Example 1

The following example configures the SPNEGO TAI to intercept HTTP requests that contain IE 6 in the user agent request header. The SPNEGO TAI uses the SPN of HTTP/myhost.ibm.com@<default_realm> to authenticate the request originator. Then the example modifies the value of the filter custom property that was defined and changes it from user-agent%=IE 6 to host==myhost.company.com.

```
$AdminTask addSpnegoTAIProperties -host myhost.ibm.com -filter user-agent%=IE 6
$AdminTask modifySpnegoTAIProperties -spnId 1 -filter host==myhost.company.com
```

Example 2

This is an example of modifying the SPNEGO TAI for SPN1 properties to add a filter for host central01.austin.ibm.com.

```
wsadmin>$AdminTask modifySpnegoTAIProperties -interactive
Modify SPNEGO TAI properties
```

Modify SPNEGO TAI configuration properties

```
*Service Principal Name identifier (spnId): 1
Host name in Service Principal Name (host): central01.austin.ibm.com
HTTP header filter rule (filter): request-url!=noSPNEGO;request-url%=snoop
Name of class used to filter HTTP requests (filterClass):
SPNEGO not supported browser response (noSpnegoPage):
NTLM Token received browser response (ntlmTokenPage):
```


Trim User Name browser response (trimUserName):

Modify SPNEGO TAI properties

F (Finish)
C (Cancel)

Select [F, C]: [F] f

```
WASX7278I: Generated command line: $AdminTask modifySpnegoTAIProperties {-spnId
1 -host w2003secdev.austin.ibm.com -filter request-url!=noSPNEGO;request-url%=sn
oop}
com.ibm.ws.security.spnego.SPN1.filter=request-url!=noSPNEGO;request-url%=snoop
com.ibm.ws.security.spnego.SPN1.hostName=central01.austin.ibm.com
wsadmin>
```

Deleting SPNEGO TAI properties using the wsadmin utility (deprecated):

You use the wsadmin utility to delete properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

About this task

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Verify that end-user desktop browsers are configured to support SPNEGO authentication, that the SPNEGO TAI is enabled, that the Java virtual machine (JVM) property is set and, that WebSphere Application Server is configured to enable the operation of the SPNEGO TAI.

You use the wsadmin utility to configure the SPNEGO TAI for WebSphere Application Server:

1. Start WebSphere Application Server.
2. Start the command-line utility by running the **wsadmin** command from the *app_server_root/bin* directory.
3. At the **wsadmin** prompt, enter the following command:

```
$AdminTask deleteSpnegoTAIProperties
```

You can use the following parameters with this command:

Option	Description
<spnId>	This is an optional parameter. It is the SPN identifier for the group of custom properties that are to be deleted with this command. If you do not specify this parameter, all SPNEGO TAI custom properties are deleted.

Results

SPNEGO TAI properties are deleted for this WebSphere Application Server.

Example

Example 1

The following example deletes all the SPNEGO TAI properties for SPN2

```
wsadmin>$AdminTask deleteSpnegoTAIProperties {-spnId 2}
```

Example 2

The following example deletes all SPNEGO TAI properties

```
wsadmin>$AdminTask deleteSpnegoTAIProperties  
com.ibm.ws.security.spnego.SPN1.filter=request-uri!=noSPNEGO;request-uri%=snoop  
com.ibm.ws.security.spnego.SPN1.hostName=central01.austin.ibm.com  
com.ibm.ws.security.spnego.SPN2.hostName=wssecpd.austin.ibm.com  
wsadmin>
```

Displaying SPNEGO TAI properties using the wsadmin utility (deprecated):

You use the wsadmin utility to display the properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

About this task

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Verify that end-user desktop browsers are configured to support SPNEGO authentication, that the SPNEGO TAI is enabled, that the Java virtual machine (JVM) property is set and, that WebSphere Application Server is configured to enable the operation of the SPNEGO TAI.

You use the wsadmin utility to configure the SPNEGO TAI for WebSphere Application Server:

1. Start WebSphere Application Server.
2. Start the command-line utility by running the **wsadmin** command from the *app_server_root/bin* directory.
3. At the **wsadmin** prompt, enter the following command:

```
$AdminTask showSpnegoTAIProperties
```

You can use the following parameters with this command:

Option	Description
<spnId>	This is an optional parameter. It is the service principal name (SPN) identifier for the group of custom properties that are to be displayed with this command. If you do not specify this parameter, all SPNEGO TAI custom properties are displayed.

Results

SPNEGO TAI properties are displayed for this WebSphere Application Server.

Example

Example 1

The following example displays all SPNEGO TAI properties.

```
wsadmin>$AdminTask showSpnegoTAIProperties
com.ibm.ws.security.spnego.SPN1.filter=request-url!=noSPNEGO;request-url%=snoop
com.ibm.ws.security.spnego.SPN1.hostName=central01.austin.ibm.com
com.ibm.ws.security.spnego.SPN2.hostName=wssecpd.austin.ibm.com
wsadmin>
```

Example 2

The following example displays SPNEGO TAI properties for SPN1 and host, central01.austin.ibm.com.

```
wsadmin>$AdminTask showSpnegoTAIProperties -interactive
Show SPNEGO TAI configuration properties.
```

Display SPNEGO TAI configuration properties.

Service Principal Name identifier (spnId): 1

Show SPNEGO TAI configuration properties.

F (Finish)

C (Cancel)

Select [F, C]: [F]

WASX7278I: Generated command line: \$AdminTask showSpnegoTAIProperties {-spnId 1}

```
com.ibm.ws.security.spnego.SPN1.filter=request-url!=noSPNEGO;request-url%=snoop
com.ibm.ws.security.spnego.SPN1.hostName=central01.austin.ibm.com
com.ibm.ws.security.spnego.SPN1.trimUserName=true
wsadmin>
```

SPNEGO TAI custom properties configuration (deprecated):

The Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) custom configuration properties control different operational aspects of the SPNEGO TAI. You can specify different property values for each application server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Each of the properties defined in the following table is specified in the Custom Properties panel for the SPNEGO TAI using the administrative console facility. For convenience, you can optionally place these properties in a properties file. In this case, the SPNEGO TAI loads the configuration properties from the file instead of the Custom Properties panel definition. Refer to `com.ibm.ws.security.spnego.propertyReloadFile` property as defined in “SPNEGO TAI JVM configuration custom properties (deprecated)” on page 344.

To assign unique property names that identify each possible SPN, an `SPN<id>` is embedded in the property name and used to group the properties that are associated with each SPN. The `SPN<id>`s are numbered sequentially for each property group.

Table 10.

Property Name	Required	Default Value
<code>com.ibm.ws.security.spnego.SPN<id>.hostName</code>	Yes	None
<code>com.ibm.ws.security.spnego.SPN<id>.filterClass</code>	No	See the description that follows.
<code>com.ibm.ws.security.spnego.SPN<id>.filter</code>	No	See the description that follows.
<code>com.ibm.ws.security.spnego.SPN<id>.enableCredDelegate</code>	No	false
<code>com.ibm.ws.security.spnego.SPN<id>.spnegoNotSupportedPage</code>	No	See the description that follows.
<code>com.ibm.ws.security.spnego.SPN<id>.NTLMTOKENReceivedPage</code>	No	See the description that follows.
<code>com.ibm.ws.security.spnego.SPN<id>.trimUserName</code>	No	true

Note: The following commands tasks can be used to operate on these SPNEGO TAI properties:

- “Adding SPNEGO TAI properties using the wsadmin utility (deprecated)” on page 329
- “Deleting SPNEGO TAI properties using the wsadmin utility (deprecated)” on page 334
- “Modifying SPNEGO TAI properties using the wsadmin utility (deprecated)” on page 331
- “Displaying SPNEGO TAI properties using the wsadmin utility (deprecated)” on page 335

Related concepts

“Single sign-on for HTTP requests using SPNEGO TAI (deprecated)” on page 320

WebSphere Application Server provides a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources in WebSphere Application Server.

Related tasks

“Mapping Kerberos client principal name to WebSphere user registry ID for SPNEGO TAI (deprecated)” on page 347

You can use a system programming interface to customize the behavior of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) by implementing arbitrary mappings of the end-user’s identity, which is retrieved from Microsoft Active Directory to the identity that is used in the WebSphere Application Server security registry.

com.ibm.ws.security.spnego.SPN<id>.hostName:

This property is required. It specifies the hostname in the SPN used by the SPNEGO TAI to establish a Kerberos secure context. It has no default value.

Note: The hostname is the long form of hostname. For example, `myHostName.austin.ibm.com`.

The Kerberos SPN is a string of the form `HTTP/hostname@realm`. The complete SPN is used with the Java Generic Security Service (JGSS) by the SPNEGO provider to obtain the security credential and security context that are used in the authentication process.

com.ibm.ws.security.spnego.SPN<id>.filterClass:

This property is optional. It specifies the name of the Java class that is used by the SPNEGO TAI to select which HTTP requests are subject to SPNEGO authentication.

If no class is specified, the default `com.ibm.ws.security.spnego.HTTPHeaderFilter` implementation class is used. The Java class that is specified must implement the `com.ibm.wsspi.security.spnego.SpnegoFilter` interface. A default implementation of this interface is provided. Specify the

`com.ibm.ws.security.spnego.HTTPHeaderFilter` class to use the default implementation. This class uses the selection rules specified with the `com.ibm.ws.security.spnego.SPN<id>.filter` property.

com.ibm.ws.security.spnego.SPN<id>.filter:

This property is optional. It defines the filtering criteria that is used by the specified class with the `com.ibm.ws.security.spnego.SPN<id>.filterClass` property. It defines arbitrary criteria that is meaningful to the implementation class used.

The `com.ibm.ws.security.spnego.HTTPHeaderFilter` default implementation class uses this property to define a list of selection rules that represent conditions that are matched against the HTTP request headers to determine whether or not the HTTP request is selected for SPNEGO authentication.

Each condition is specified with a key-value pair, separated from each other by a semicolon. The conditions are evaluated from left to right, as they display in the specified property. If all conditions are met, the HTTP request is selected for SPNEGO authentication.

The key and value in the key-value pair are separated by an operator that defines which condition is checked. The key identifies an HTTP request header to extract from the request and its value is compared with the value that is specified in the key-value pair according to the operator specification. If the header that is identified by the key is not present in the HTTP request, the condition is treated as not being met.

Any of the standard HTTP request headers can be used as the key in the key-value pairs. Refer to the HTTP specification for the list of valid headers. In addition, two keys are defined to extract information from the request, also useful as a selection criterion, which is not available through standard HTTP request headers. The `remote-address` key is used as a pseudo header to retrieve the remote TCP/IP address of the client application that sent the HTTP request. The `request-URL` key is used as a pseudo header to retrieve the URL that is used by the client application to make the request. The interceptor uses the result of the `getRequestURL` operation in the `javax.servlet.http.HttpServletRequest` interface to construct the Web address. If a query string is present, the result of the `getQueryString` operation in the same interface is also used. In this case, the complete URL is constructed as follows:

```
String url = request.getRequestURL() + '?' + request.getQueryString();
```

The following operators and conditions are defined:

Table 11. Filter conditions and operations

Condition	Operator	Example
Match exactly	= = Arguments are compared as equal.	host=host.my.company.com
Match partially (includes)	%= Arguments are compared with a partial match being valid.	user-agent%=IE 6
Match partially (includes one of many)	^= Arguments are compared with a partial match being valid for one of many arguments specified.	request-url^=webApp1 webApp2 webApp3
Does not match	!= Arguments are compared as not equal.	request-url!=noSPNEGO

Table 11. Filter conditions and operations (continued)

Condition	Operator	Example
Greater than	> Arguments are compared lexographically as greater than.	remote-address>192.168.255.130
Less than	< Arguments are compared lexographically as less than.	remote-address<192.168.255.135

com.ibm.ws.security.spnego.SPN<id>.enableCredDelegate:

This property is optional. It indicates whether or not the Kerberos delegated credentials are stored by the SPNEGO TAI. This property enables the capability for an application to retrieve the stored credentials and propagate them to other applications downstream for additional SPNEGO authentication.

This property requires use of the advanced Kerberos credential delegation feature and requires development of custom logic by the application developer. The developer must interact directly with the Kerberos Ticket Granting Service (TGS) to obtain a Ticket Granting Ticket (TGT) using the delegated Kerberos credentials on behalf of the end-user who originated the request. The developer must also construct the appropriate Kerberos SPNEGO token and include it in the HTTP request to continue the downstream SPNEGO authentication process, including handling additional SPNEGO challenge-response exchange, if necessary.

com.ibm.ws.security.spnego.SPN<id>.spnegoNotSupportedPage:

This property is optional. It specifies the Web address of a resource that contains the content that the SPNEGO TAI includes in the HTTP response that the (browser) client application displays if it does not support SPNEGO authentication. It can specify a Web (http://) or a file (file://) resource.

If this property is not specified or the interceptor cannot find the specified resource, the following content is used:

```
<html><head><title>SPNEGO authentication is not supported</title></head>
<body>SPNEGO authentication is not supported on this client</body></html>;
```

com.ibm.ws.security.spnego.SPN<id>.NTLMTokenReceivedPage:

This property is optional. It specifies the Web address of a resource that contains the content that the SPNEGO TAI includes in the HTTP response that the (browser) client application displays when the SPNEGO token is received by the interceptor when the challenge-response handshake contains a NT LAN Manager (NTLM) token instead of the expected SPNEGO token.

It can specify a Web (http://) or a file (file://) resource. If this property is not specified or the interceptor cannot find the specified resource, the following content is used:

```
<html><head><title>An NTLM Token was received.</title></head>
<body>Your browser configuration is correct, but you have not logged into a supported
Microsoft(R) Windows(R) Domain.
<p>Please login to the application using the normal login page.</html>
```

com.ibm.ws.security.spnego.SPN<id>.trimUserName:

This property is optional. It specifies whether (true) or not (false) the SPNEGO TAI is to remove the suffix of the principal user name, starting from the "@" that precedes the Kerberos realm name.

If this property is set to `true`, the suffix of the principal user name is removed. If this property is set to `false`, the suffix of the principal name is retained. The default value used is `true`. For example,

When `com.ibm.ws.security.spnego.SPN<id>.trimUserName = true`

`bobsmith@myKerberosRealm` becomes `bobsmith`

When `com.ibm.ws.security.spnego.SPN<id>.trimUserName = false`

`bobsmith@myKerberosRealm` remains `bobsmith@myKerberosRealm`

SPNEGO TAI configuration requirements (deprecated):

The configuration that is used by the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) on each selected application server is governed by various system requirements.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

The following list of configuration requirements highlights those attributes, properties, qualities, restrictions, exclusions, inclusions, and dependencies that you need to be aware of when planning a WebSphere Application Server configuration that incorporates the use of the SPNEGO TAI.

Table 12. SPNEGO TAI requirements

Function item	Description
SPNEGO TAI	The SPNEGO TAI is a server side solution in WebSphere Application Server. Client-side applications are responsible for generating the SPNEGO token for use by the SPNEGO TAI.
Microsoft Windows	Windows 2000 or Windows 2003 servers with Active Directory domain and its associated Kerberos key distribution center (KDC) is required.
Client application (browser or .NET client)	A browser (client application) or .NET client that supports the SPNEGO authentication mechanism, as defined in IETF RFC 2478 is required.
Simple and Protected GSS-API Negotiation Mechanism (SPNEGO)	SPNEGO authentication, as defined in IETF RFC 2478 is used.
Internet browsers	<ul style="list-style-type: none"> • Use Microsoft Internet Explorer version 5.5 or higher • Use Mozilla Firefox version 1.0
Kerberos Level	Kerberos version 5 is required.
WebSphere Application Server	Version 7.0 is required.
Java SDK level	Java 6.0 SDK is required.
Encryption Types	RC4-HMAC encryption is only supported when using a Windows 2003 Server as Kerberos key distribution center (KDC) and is not supported with a Windows 2000 Server.
J2EE client	Client application (browser or .NET client) A browser (client application) or .NET client that supports the SPNEGO authentication mechanism, as defined in IETF RFC 2478 is required.

Configuring the client browser to use SPNEGO TAI (deprecated)

You can configure your browser to utilize the Simple and Protected GSS-API Negotiation (SPNEGO) mechanism. Authentication of your browser requests are processed by the SPNEGO trust association interceptor (TAI) in the WebSphere Application Server.

Before you begin

You need to know how to display and set options in the Microsoft Internet Explorer browser or any other browser (such as Firefox). You must have a browser installed that supports SPNEGO authentication.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

Complete the following steps to ensure that your Microsoft Internet Explorer browser is enabled to perform SPNEGO authentication.

1. At the desktop, log in to the windows active directory domain.
2. Activate Internet Explorer.
3. In the Internet Explorer window, click **Tools > Internet Options > Security** tab.
4. Select the **Local intranet** icon and click **Sites**.
5. In the Local intranet window, ensure that the "check box" to include all local (intranet) not listed in other zones is selected, then click **Advanced**.
6. In the **Local intranet** window, fill in the Add this Web site to the zone field with the Web address of the host name so that the single sign-on (SSO) can be enabled to the list Web sites shown in the Web sites field. Your site information technology staff provides this information. Click **OK** to complete this step and close the Local intranet window.
7. On the **Internet Options** window, click the **Advanced** tab and scroll to **Security settings**. Ensure that the **Enable Integrated Windows Authentication (requires restart)** box is selected.
8. Click **OK**. Restart your Microsoft Internet Explorer to activate this configuration.

Results

Complete the following steps to ensure that your Firefox browser is enabled to perform SPNEGO authentication.

1. At the desktop, log in to the windows active directory domain.
2. Activate Firefox.
3. At the address field, type **about:config**.
4. In the Filter, type **network.n**
5. Double click on **network.negotiate-auth.trusted-uris**. This preference lists the sites that are permitted to engage in SPNEGO Authentication with the browser. Enter a comma-delimited list of trusted domains or URLs.

Note: You must set the value for **network.negotiate-auth.trusted-uris**.

6. If the deployed SPNEGO solution is using the advanced Kerberos feature of Credential Delegation double click on **network.negotiate-auth.delegation-uris**. This preference lists the sites for which the browser may delegate user authorization to the server. Enter a comma-delimited list of trusted domains or URLs.
7. Click **OK**. The configuration appears as updated.
8. Restart your Firefox browser to activate this configuration.

Your Internet browser is properly configured for SPNEGO authentication. You can use applications that are deployed in WebSphere Application Server that use secured resources without being repeatedly requested for an ID and password.

Configuring JVM custom properties, filtering HTTP requests, and enabling SPNEGO TAI in WebSphere Application Server (deprecated)

Performing this task helps you, as Web administrator, to ensure that WebSphere Application Server is configured to enable the operation of the Simple and Protected GSS-API Negotiation mechanism (SPNEGO) trust association interceptor (TAI) with the required Java virtual machine (JVM) property and with the appropriate filtering of HTTP requests.

Before you begin

You need to know how to use the WebSphere Application Server administrative console to manage the security configuration and have the proper authority to modify the security configuration of the application server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

Verify the configuration of your SPNEGO TAI. The deployment of the SPNEGO TAI can vary from a single WebSphere Application Server system on which a single application is running to a large multinode WebSphere Application Server Network Deployment (ND) cell, with dozens of application servers, hosting many applications. Every SPNEGO TAI is installed at the cell level. You must be aware of your particular SPNEGO TAI configuration.

The default behavior of the SPNEGO TAI is to not intercept HTTP requests. This default behavior ensures that the SPNEGO TAI can be installed into an existing cell, configured for a single application server and not change any other application servers in the cell. Other WebSphere Application Servers can run exactly as before within a given configuration.

Decide whether or not to use the sample `SPN<id>.filterClass` and determine the exact filter properties to use.

Note: The default behavior of the SPNEGO TAI is to use the `com.ibm.ws.security.spnego.SPN<id>.filterClass` and intercept all requests.

If the default behavior is not appropriate, you can use a customer provided class, or extend or modify the sample class as required. The system programmer interface, `com.ibm.ws.security.spnego.SpnegoFilter` allows you to implement a custom filter to determine whether or not to intercept a particular HTTP request. With the default implementation, you can set filter rules for coarse as well as fine-grained criteria in selecting which HTTP requests to intercept.

Note: For an alternative to the steps below for enabling the SPNEGO TAI, you can use scripting to perform the operation. See “Enabling the SPNEGO TAI as JVM custom property using scripting (deprecated)” for the details.

Complete the following steps to enable the operation of the SPNEGO TAI with your selected filtering and with the JVM required property.

1. Log on to WebSphere Application Server administrative console.
2. Click **Servers > Application servers**.
3. Select the appropriate server. Under Server Infrastructure, expand **Java and process management > Process Definition**.
4. Click **Java virtual machine**. Under Additional Properties, click **Custom Properties**. Create a new custom property, if required, by clicking **New**, then code `com.ibm.ws.security.spnego.isEnabled` in the name field and `true` in the value field.
5. Click **Apply > OK** to save the configuration
6. Identify when the SPNEGO TAI intercepts a given request. A set of filter properties is provided, but you must determine what is appropriate and modify the `com.ibm.ws.security.spnego.SPN<id>.filterClass` accordingly.

Results

The application server is configured and ready to provide a single sign-on environment for end users who have successfully authenticated in a Microsoft Active Directory domain. You must restart each application server that is configured for SPNEGO Web authentication. Then your SPNEGO TAI is set to filter HTTP request when it is operating.

Enabling the SPNEGO TAI as JVM custom property using scripting (deprecated):

You use the `wsadmin` utility to enable the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Before you begin

Before starting this task, the `wsadmin` tool must be running. See the Starting the `wsadmin` scripting client article for more information.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

Perform the following steps to enable the SPNEGO TAI:

1. Identify the server and assign it to the `server1` variable:

- Using Jacl:

```
set server1 [AdminConfig getid /Cell:mycell/Node:mynode/Server:server1/]
```

- Using Jython:

```
server1 = AdminConfig.getid("/Cell:mycell/Node:mynode/Server:server1/")
print server1
```

Example output:

```
server1(cells/mycell/nodes/mynode|servers/seerver1|server.xml#Server_1)
```

2. Identify the Java virtual machine (JVM) belonging to this server and assign it to the `jvm` variable:

- Using Jacl:

```
set jvm [$AdminConfig list JavaVirtualMachine $server1]
```

- Using Jython:

```
jvm = AdminConfig.list('JavaVirtualMachine',server1)
```

Example output:

```
(cells/mycell/nodes/mynode/servers/server1:server.xml#JavaVirtualMachine_1)  
(cells/mycell/nodes/mynode/servers/server1:server.xml#JavaVirtualMachine_2)
```

3. Identify the controller JVM of the server:

- Using Jacl:

```
set cjvm [lindex $jvm 0]
```

- Using Jython:

```
# get line separator  
import java  
lineSeparator = java.lang.System.getProperty('line.separator')  
arrayJVMs = jvm.split(lineSeparator)  
cjvm = arrayJVMs[0]
```

4. Modify the generic JVM arguments to enable SPNEGO TAI:

- Using Jacl:

```
set attr_name          [list name com.ibm.ws.security.spnego.isEnabled]  
set attr_value         [list value true]  
set attr_required     [list required false]  
set attr_description  [list description "Enabled SPNEGO TAI"]
```

```
set attrs [list $attr_name $attr_value $attr_required $attr_description]
```

```
$AdminConfig create Property $cjvm $attrs
```

- Using Jython:

```
attr_name = ['name', "com.ibm.ws.security.spnego.isEnabled"]  
attr_value = ['value', "true"]  
attr_required = ['required', "false"]  
attr_description = ['description', "Enabled SPNEGO TAI"]  
attr_list = [attr_name, attr_value, attr_required, attr_description]  
property=['systemProperties',[attr_list]]  
AdminConfig.modify(cjvm, [property])
```

5. Save the configuration changes. See the Saving configuration changes with the `wsadmin` tool article for more information.
6. In a network deployment environment only, synchronize the node. See the Synchronizing nodes with the `wsadmin` tool article for more information.

SPNEGO TAI JVM configuration custom properties (deprecated):

Java virtual machine (JVM) custom properties control the operation of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI).

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

The following JVM custom properties control operation of the SPNEGO TAI. Different custom property values can be specified for each application server.

Table 13. JVM configuration custom properties

Custom Property Name	Required	Value Type	Default Value	Recommended Value
com.ibm.ws.security.spnego.isEnabled	No	Boolean	False	True
com.ibm.ws.security.spnego.propertyReloadFile	No	String	None	For Windows C:\temp\TAI.props For UNIX /tmp/TestTAI.Properties
com.ibm.ws.security.spnego.propertyReloadTimeout	No	Integer	None	120

com.ibm.ws.security.spnego.isEnabled

Use this custom property to enable or disable operation of the SPNEGO TAI in a given application server. When set to `false`, the SPNEGO TAI is disabled and not used by the Web authentication module for authenticating any Web requests. When set to `true`, the SPNEGO TAI is enabled and used by the Web authentication module for authenticating any Web requests.

com.ibm.ws.security.spnego.propertyReloadFile

Use this custom property to identify the file that contains configuration properties for the SPNEGO TAI, when it is not convenient to stop and restart the application server. The properties contained in this file can be reloaded to configure the SPNEGO TAI.

Note: The properties that are defined in the specified file override any properties defined using the administrative console.

A sample of this reload file follows:

```
#####
# Template properties files for SPNEGO TAI
#
# Where possible defaults have been provided.
#
#####

#-----
# Hostname
#-----
#com.ibm.ws.spnego.SPN1.HostName=wsecurity.austin.ibm.com

#-----
# (Optional) SpnegoNotSupportedPage
#-----
#com.ibm.ws.spnego.SPN1.SpnegoNotSupportedPage=

#-----
# (Optional) NTLMTokenReceivedPage
#-----
#com.ibm.ws.spnego.SPN1.NTLMTokenReceivedPage=

#-----
# (Optional) FilterClass
#-----
#com.ibm.ws.spnego.SPN1.FilterClass=com.ibm.ws.spnego.HTTPHeaderFilter

#-----
# (Optional) Filter
#-----
#com.ibm.ws.spnego.SPN1.Filter=
```

Note: If `com.ibm.ws.security.spnego.propertyReloadFile` custom property is set, but the `com.ibm.ws.security.spnego.propertyReloadTimeout` custom property is not, then the SPNEGO TAI is not initialized.

com.ibm.ws.security.spnego.propertyReloadTimeout

Use this custom property to specify a time interval in seconds that elapses after which the SPNEGO TAI reloads the configuration properties. Also, the SPNEGO TAI reloads the configuration properties if the file that is identified by the `com.ibm.ws.security.spnego.propertyReloadFile` custom property changed since the last time the configuration custom properties were retrieved. This time interval in seconds must be specified as a positive integer.

Note:

- If the `com.ibm.ws.security.spnego.propertyReloadFile` custom property and the `com.ibm.ws.security.spnego.propertyReloadTimeout` custom property are not set, then the SPNEGO TAI properties are only loaded once from the SPNEGO TAI custom properties defined in the WebSphere Application Server configuration data. This one time loading occurs when the JVM is initialized.
- If `com.ibm.ws.security.spnego.propertyReloadTimeout` custom property is set, but the `com.ibm.ws.security.spnego.propertyReloadFile` custom property is not, then the SPNEGO TAI is not initialized. “Configuring JVM custom properties, filtering HTTP requests, and enabling SPNEGO TAI in WebSphere Application Server (deprecated)” on page 342 or how to configure the JVM custom properties for SPNEGO TAI.

Note: You can also use the `wsadmin` command for the `AdminConfig` scripting object to interactively set the `com.ibm.ws.security.spnego.isEnabled` custom property. See “Enabling the SPNEGO TAI as JVM custom property using scripting (deprecated)” on page 343 for more information.

The following custom properties are not used directly by the SPNEGO TAI; however, they affect the operation of the core security runtime and can also be used for problem determination.

Table 14. JVM configuration custom properties

Custom Property Name	Required	Value Type	Default Value	Recommended Value
<code>java.security.properties</code>	No	String	None	
<code>com.ibm.security.jgss.debug</code>	No	String	None	"off" or "all"
<code>com.ibm.security.krb5.Krb5Debug</code>	No	String	None	"off" or "all"
<code>javax.security.auth.useSubjectCredsOnly</code>	Yes	Boolean	True	False

java.security.properties

This property is optional. It can be used when different application servers in a cell have different security requirements and it is not convenient to modify the global `java.security` file for the entire cell. In such situations, the `java.security.properties` custom property is used to specify the location of the `java.security` file used by the JVM for each application server.

com.ibm.security.jgss.debug

This custom property is optional. It can be used to collect diagnostic trace information for problem determination in the Java Generic Security Service (JGSS) application programmer interface (API) implementation. The value can be set to `all` or `off` to enable or disable tracing, respectively. See Java Generic Security Service User’s Guide for specific JGSS API information.

com.ibm.security.krb5.Krb5Debug

This custom property is optional. It can be used to collect additional diagnostic trace information for problem determination in the JGSS implementation. The value can be set to `all` or `off` to enable or disable tracing, respectively.

javax.security.auth.useSubjectCredsOnly

JGSS includes an optional Java Authentication and Authorization Service (JAAS) login facility that

saves Principal credentials and secret keys in the Subject of the application's JAAS login context. JGSS retrieves credentials and secret keys from the Subject by default. This feature can be disabled by setting the Java property `javax.security.auth.useSubjectCredsOnly` to `false`.

Note: The SPNEGO TAI does not use the optional JAAS login module. The `javax.security.auth.useSubjectCredsOnly` property must be set to `false`.

Mapping Kerberos client principal name to WebSphere user registry ID for SPNEGO TAI (deprecated)

You can use a system programming interface to customize the behavior of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) by implementing arbitrary mappings of the end-user's identity, which is retrieved from Microsoft Active Directory to the identity that is used in the WebSphere Application Server security registry.

Before you begin

You need to perform some administrative tasks in the WebSphere Application Server environment to use SPNEGO TAI and to ensure that the requester's identity matches the identity in the WebSphere Application Server user registry.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Note: Make sure the following tasks have been performed successfully:

1. Configuring the Web browser to use SPNEGO. See "Configuring the client browser to use SPNEGO TAI (deprecated)" on page 341
2. Configuring Java virtual machine (JVM) properties, custom SPNEGO TAI properties, and enabling the SPNEGO TAI. See "Configuring JVM custom properties, filtering HTTP requests, and enabling SPNEGO TAI in WebSphere Application Server (deprecated)" on page 342

About this task

In the simplest deployment of the SPNEGO TAI, it is assumed that the requester's identity in the WebSphere Application Server user registry is identical to the identity retrieved. This is the case when Microsoft Windows Active Directory server is the lightweight directory access protocol (LDAP) server used in WebSphere Application Server. This is default behavior of the SPNEGO TAI.

You do not need to use this simple deployment of the SPNEGO TAI. WebSphere Application Server can use a different registry, such as a local OS, LDAP, or custom registry instead of the Microsoft Active Directory. If WebSphere Application Server uses a different registry than the Microsoft Active Directory, then a mapping from the Microsoft Windows user Id to a WebSphere Application Server user Id is necessary.

Use the JAAS custom login module to perform any custom mapping of a client Kerberos principal name from the Microsoft Active Directory to the WebSphere user registry identity. The JAAS custom login module is a plug-in mechanism that is defined for authenticating incoming and outgoing requests in WebSphere Application Server and is inserted before the `ltpaLoginModule`. The JAAS custom login module retrieves a client Kerberos principal name in the `javax.security.auth.Subject` using `subject.getPrincipals(KerberosPrincipal.class)` method, maps the client Kerberos principal name to

the WebSphere user registry identity, and inserts the mapping identity in the hash table property `com.ibm.wsspi.security.cred.userId`. The `ltpaLoginModule` then uses the mapped identity to create a `WSCredential`.

Note: The custom login module can also supply the full set of security properties in the `javax.security.auth.Subject` in the `com.ibm.wsspi.security.tai.TAIResult` to fully assert the mapped identity. When the identity is fully asserted, the `wsMapDefaultInboundLoginModule` maps those security properties to a `WSCredential`.

A sample of the custom login module follows:

```
package com.ibm.ws.security.server.lm;

import java.util.Map;
import java.lang.reflect.Array;
import javax.security.auth.Subject;
import javax.security.auth.callback.*;
import javax.security.auth.login.LoginException;
import javax.security.auth.spi.LoginModule;
import javax.security.auth.kerberos.*;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.AttributeNameConstants;

/**
 *
 * @author IBM Corporation
 * @version 1.0
 * @since 1.0
 */

public class sampleSpnegoMappingLoginModule implements LoginModule {
    /**
     *
     * Constant that represents the name of this mapping module. Whenever this sample
     * code is used to create a class with a different name, this value should be changed.
     */
    private final static String MAPPING_MODULE_NAME = "com.ibm.websphere.security.sampleSpnegoMappingLoginModule";

    private String mapUid = null;
    /**
     * Construct an uninitialized WSLoginModuleImpl object.
     */
    public sampleSpnegoMappingLoginModule() {
        debugOut("sampleSpnegoMappingLoginModule() entry");
        debugOut("sampleSpnegoMappingLoginModule() exit");
    }

    /**
     * Initialize this login module.
     *
     * This is called by the LoginContext after this login module is
     * instantiated. The relevant information is passed from the LoginContext
     * to this login module. If the login module does not understand any of the data
     * stored in the sharedState and options parameters,
     * they can be ignored.
     *
     * @param subject The subject to be authenticated.
     * @param callbackHandler
     *             A CallbackHandler for communicating with the end user to gather
     *             login information (e.g., username and password).
     * @param sharedState
     *             The state shared with other configured login modules.
     * @param options The options specified in the login configuration for this particular login module.
     */
}
```

```

*/
public void initialize(Subject subject, CallbackHandler callbackHandler,
    Map sharedState, Map options) {
    debugOut("initialize(subject = \"" + subject.toString() +
        "\", callbackHandler = \"" + callbackHandler.toString() +
        "\", sharedState = \"" + sharedState.toString() +
        "\", options = \"" + options.toString() + "\")");

    this.subject = subject;
    this.callbackHandler = callbackHandler;
    this.sharedState = sharedState;
    this.options = options;

    debug = "true".equalsIgnoreCase((String)this.options.get("debug"));

    debugOut("initialize() exit");
}

/**
 *
 * Method to authenticate a Subject (phase 1).
 *
 *
 *
 * This method authenticates a Subject. It uses CallbackHandler to gather
 * the Subject information, like username and password for example, and verify these
 * information. The result of the authentication is saved in the private state within
 * this login module.
 *
 *
 * @return true if the authentication succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the authentication fails.
 */
public boolean login() throws LoginException
{
    debugOut("sampleSpnegoMappingLoginModule.login() entry");

    boolean succeeded = false;
    java.util.Set krb5Principals= subject.getPrincipals(KerberosPrincipal.class);
    java.util.Iterator krb5PrincIter = krb5Principals.iterator();

    while (krb5PrincIter.hasNext()) {
        Object princObj = krb5PrincIter.next();
        debugOut("Kerberos principal name: "+ princObj.toString());

        if (princObj != null && princObj.toString().equals("utle@WSSEC.AUSTIN.IBM.COM")){
            mapUid = "user1";
            debugOut("mapUid: "+mapUid);

            java.util.Hashtable customProperties = (java.util.Hashtable)
                sharedState.get(AttributeNameConstants.WSCREDENTIAL_PROPERTIES_KEY);
            if (customProperties == null) {
                customProperties = new java.util.Hashtable();
            }
            succeeded = true;
            customProperties.put(AttributeNameConstants.WSCREDENTIAL_USERID, mapUid);

            Map<String,java.util.Hashtable>
            mySharedState=(Map<String,java.util.Hashtable>)sharedState;
            mySharedState.put((AttributeNameConstants.WSCREDENTIAL_PROPERTIES_KEY.customProperties)

            debugOut("Add a mapping user ID to Hashtable, mapping ID = "+mapUid);

        }
        debugOut("login() custom properties = " + customProperties);
    }
}

```

```

    }
}

succeeded = true;
debugOut("sampleSpnegoMappingLoginModule.login() exit");

return succeeded;
}

/**
 *
 * Method to commit the authentication result (phase 2).
 *
 *
 * This method is called if the LoginContext's overall authentication
 * succeeded (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL login module
 * succeeded).
 *
 *
 * @return true if the commit succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the commit fails.
 */
public boolean commit() throws LoginException
{
    debugOut("commit()");

    debugOut("commit()");

    return true;
}

/**
 * Method to abort the authentication process (phase 2).
 *
 *
 * This method is called if the LoginContext's overall authentication
 * failed (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL login module
 * did not succeed).
 *
 *
 *
 * If this login module's authentication attempt succeeded, then this method cleans
 * up the previous state saved in phase 1.
 *
 *
 * @return true if the abort succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the abort fails.
 */
public boolean abort() throws LoginException {
    debugOut("abort() entry");
    debugOut("abort() exit");
    return true;
}

/**
 * Method which logs out a Subject.
 *
 *
 * @return true if the logout succeeded, or false
 *         if this login module should be ignored.
 * @exception LoginException
 *         If the logout fails.
 */

```

```

public boolean logout() throws LoginException
{
    debugOut("logout() entry");
    debugOut("logout() exit");

    return true;
}

private void cleanup()
{
    debugOut("cleanup() entry");
    debugOut("cleanup() exit");
}

/*
 *
 * Private method to print trace information. This implementation uses System.out
 * to print trace information to standard output, but a custom tracing system can
 * be implemented here as well.
 */
private void debugOut(Object o)
{
    System.out.println("Debug: " + MAPPING_MODULE_NAME);
    if (o != null) {
        if (o.getClass().isArray()) {
            int length = Array.getLength(o);
            for (int i = 0; i < length; i++) {
                System.out.println("\t" + Array.get(o, i));
            }
        } else {
            System.out.println("\t" + o);
        }
    }
}

private Subject subject;
private CallbackHandler callbackHandler;
private Map sharedState;
private Map options;

protected boolean debug = false;
}

```

Results

Using the custom login module, Microsoft Active Directory identities are mapped to the WebSphere Application Server's security registry and the behavior of the SPNEGO TAI is customized.

Single sign-on capability with SPNEGO TAI - checklist (deprecated)

WebSphere Application Server provides a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources in WebSphere Application Server. To deploy and use the SPNEGO TAI you need to examine your installation and decide on how best to configure the SPNEGO TAI.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

Lightweight Third Party Authentication (LTPA) is the default authentication mechanism for WebSphere Application Server. However, you may need to configure LTPA prior to configuring the SPNEGO TAI. LTPA is the required authentication mechanism for all trust association interceptors. You can configure LTPA by clicking **Security > Global security > Authentication mechanisms and expiration**.

Note: Enabling Web security single sign-on (SSO) is optional when you configure the SPNEGO TAI. For more information, see “Implementing single sign-on to minimize Web user authentications” on page 297.

Answer the following questions to establish how the SPNEGO TAI is deployed.

1. What is your criteria for intercepting HTTP requests?

You must decide if the SPNEGO TAI deployment will use the HTTPHeaderFilter class as the default. If you do use this class, then you must specify the exact filter properties for this class. The default behavior of the SPNEGO TAI is to use the com.ibm.ws.spnego.HTTPHeaderFilter class to intercept all requests.

If you do not use the sample com.ibm.ws.spnego.HTTPHeaderFilter class, then you must define a new class that implements the com.ibm.wsspi.security.spnego.SpnegoTAIFilter interface.

You can decide to further control what HTTP requests are intercepted using the Service Provider Programming Interface (SPI), “Filtering HTTP requests for SPNEGO TAI (deprecated)” on page 353

See “SPNEGO TAI custom properties configuration (deprecated)” on page 336 for descriptions of

- com.ibm.ws.security.spnego.SPN<id>.filterClass
- com.ibm.ws.security.spnego.SPN<id>.filter

2. Is user Id mapping to be used? If not, why not?

WebSphere Application Server enables you to define or develop a custom login module to map user IDs. See “Mapping Kerberos client principal name to WebSphere user registry ID for SPNEGO TAI (deprecated)” on page 347 for more detail about performing this mapping.

You must decide, before deploying the TAI, whether or not to use this custom login module to perform the SPNEGO TAI identity mapping

3. What type of encryption is to be used to process the SPNEGO tokens?

Microsoft Windows Active Directory supports two different Kerberos encryption types: RC4-HMAC and DES-CBC-MD5. The IBM Java Generic Security Service (JGSS) library (and SPNEGO library) support both of these encryption types.

Note: RC4-HMAC encryption is only supported with a Windows 2003 Server key distribution center (KDC). RC4-HMAC encryption is not supported when using a Windows 2000 Server as a Kerberos KDC.

4. How will you handle credential delegation?

Kerberos supports the delegation of credentials. A server that receives Kerberos credentials from a client can impersonate that client to other servers by using delegated credentials. Since SPNEGO TAI tokens are a wrapping of a Kerberos credential, a server that receives Kerberos credentials within an SPNEGO token can use those Kerberos credentials to impersonate the original user. That server can interact using SPNEGO over HTTP as a SPNEGO client to other SPNEGO servers by composing an appropriate HTTP Authorization header.

5. Will the SPNEGO TAI be deployed in a single or multiple domain name service (DNS) domain environment?

Web browsers running on Windows are sensitive to DNS domains. They only send a SPNEGO token when the target host name identifies a host name defined in the DNS domain of the client machine. You can use HTTP redirection to support this configuration with the creation of a pseudo Kerberos service principal name (SPN) in each DNS domain. All SPNs that WebSphere Application Server supports must have their secret keys available in Kerberos keytab files. To enable single sign-on

across multiple DNS domains, a separate Kerberos keytab file is generated for each SPN per domain. These individual Kerberos keytab files must be merged before they can be used by WebSphere Application Server.

6. How frequently will application servers reload the SPNEGO TAI properties?

The SPNEGO TAI has an optional property reload feature that allows the reloading of the TAI properties without restarting the Java virtual machine (JVM). This reload feature is controlled by the system properties `com.ibm.ws.security.spnego.propertyReloadFile` and `com.ibm.ws.security.spnego.propertyReloadTimeout`. These properties taken together enable the SPNEGO TAI internal properties to be reloaded from a file on the file system after a certain time period. If the `com.ibm.ws.security.spnego.propertyReloadTimeout` attribute is set to a valid integer value, and the `com.ibm.ws.security.spnego.propertyReloadFile` attribute points to a file on the file system, then each JVM reloads the SPNEGO TAI properties from the file after the timeout period expires. Also, the SPNEGO TAI properties are reloaded only if the date on the file has changed. If these reload properties are not set, then the SPNEGO TAI properties are only loaded once, at JVM initialization, from the SPNEGO TAI custom properties that are defined in WebSphere Application Server configuration data. See “SPNEGO TAI JVM configuration custom properties (deprecated)” on page 344 for more information about these reload properties.

The Windows Active Directory (Web) administrator, the WebSphere Application Server administrator, and the application team review and answer these questions to determine the best deployment and configuration settings for the SPNEGO TAI.

Filtering HTTP requests for SPNEGO TAI (deprecated)

You can use a system programming interface to customize the behavior of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) by specifying whether or not a particular HTTP request should be intercepted.

Before you begin

Before you begin, you need to understand the deployment of the SPNEGO TAI in your installation.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

About this task

Verify the configuration of your SPNEGO TAI. The deployment of the SPNEGO TAI can vary from a single WebSphere Application Server system on which a single application is running to a large multinode WebSphere Application Server Network Deployment (ND) cell, with dozens of application servers, hosting many applications. Every SPNEGO TAI is installed at the cell level. You must be aware of your particular SPNEGO TAI configuration.

The default behavior of the SPNEGO TAI is to not intercept HTTP requests. This default behavior ensures that the SPNEGO TAI can be installed into an existing cell, configured for a single application server and not change any other application servers in the cell. Other WebSphere Application Servers can run exactly as before within a given configuration.

Then decide whether or not to use the sample `SPN<id>.filter` class and determine the exact filter properties to use.

Note: The default behavior of the SPNEGO TAI is to use the `com.ibm.ws.security.spnego.SPN<id>.filter` class and intercept all requests.

If the default behavior is not appropriate, you can use a customer provided class, or extend or modify the sample class as required. The system programmer interface, `com.ibm.ws.security.spnego.SpnegoFilter` allows you to implement a custom filter to determine whether or not to intercept a particular HTTP request. With the default implementation, you can set filter rules for coarse as well as fine-grained criteria in selecting which HTTP requests to intercept.

1. Set the `com.ibm.ws.security.spnego.isEnabled` Java virtual machine (JVM) custom property to true to enable the SPNEGO TAI on any JVM.
2. Identify when the SPNEGO TAI intercepts a given request. A set of filter properties is provided, but you must determine what is appropriate and modify the `com.ibm.ws.security.spnego.SPN<id>.filter` class accordingly.

Results

Your SPNEGO TAI is set to filter HTTP requests when it is operating.

Configuring single sign-on capability with Tivoli Access Manager or WebSEAL

Use the following information to enable single sign-on to WebSphere Application Server using either WebSEAL or the plug-in for Web servers.

About this task

Either Tivoli Access Manager WebSEAL or Tivoli Access Manager plug-in for Web servers can be used as reverse proxy servers to provide access management and single sign-on (SSO) capability to WebSphere Application Server resources. With such an architecture, either WebSEAL or the plug-in authenticates users and forwards the collected credentials to WebSphere Application Server in the form of an IV Header. Two types of single sign-on are available, the TAI interface and the TAI++ interface, so named as both use WebSphere Application Server trust association interceptors (TAI). With the TAI, the end-user name is extracted from the HTTP header and forwarded to embedded Tivoli Access Manager where the end-user name is used to construct the client credential information and authorize the user. With the TAI++, all of the user credential information is available in the HTTP header and not just the user name. The TAI++ is the more efficient of the two solutions because a Lightweight Directory Access Protocol (LDAP) call is not required. TAI functionality is retained for backwards compatibility.

Complete the following tasks to enable single sign-on to WebSphere Application Server using either WebSEAL or the plug-in for Web servers. These tasks assume that embedded Tivoli Access Manager is configured for use.

1. Create a trusted user account for Tivoli Access Manager in the shared Lightweight Directory Access Protocol (LDAP) user registry. For more information, see “Creating a trusted user account in Tivoli Access Manager” on page 360.
2. Configure either WebSEAL or the Tivoli Access Manager plug-in for Web servers to work with WebSphere Application Server. For more information, see either of the following articles:
 - “Configuring WebSEAL for use with WebSphere Application Server” on page 361
 - “Configuring Tivoli Access Manager plug-in for Web servers for use with WebSphere Application Server” on page 362
3. Configure single sign-on using either the TAI or TAI++ interface. For more information, see either of the following articles:
 - “Configuring single sign-on using trust association” on page 363
 - “Configuring single sign-on using trust association interceptor ++” on page 363

Single sign-on settings

Use this page to set the configuration values for single sign-on (SSO).

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Web and SIP security > Single sign-on (SSO)**.

Enabled:

Specifies that the single sign-on function is enabled.

Web applications that use J2EE FormLogin style login pages, such as the administrative console, require single sign-on (SSO) enablement. Only disable SSO for certain advanced configurations where LTPA SSO-type cookies are not required.

Data type:	Boolean
Default:	Enabled
Range:	Enabled or Disabled

Requires SSL:

Specifies that the single sign-on function is enabled only when requests are made over HTTPS Secure Sockets Layer (SSL) connections.

Data type:	Boolean
Default:	Disable
Range:	Enable or Disable

Domain name:

Specifies the domain name (.ibm.com, for example) for all single sign-on hosts.

The application server uses all the information after the first period, from left to right, for the domain names. If this field is not defined, the Web browser defaults the domain name to the host name where the Web application is running. Also, single sign-on is then restricted to the application server host name and does not work with other application server host names in the domain.

You can specify multiple domains separated by a semicolon (;), a space (), a comma (,), or a pipe (|). Each domain is compared with the host name of the HTTP request until the first match is located. For example, if you specify `ibm.com;austin.ibm.com` and a match is found in the `ibm.com` domain first, the application server does not match the `austin.ibm.com` domain. However, if a match is not found in either `ibm.com` or `austin.ibm.com`, then the application server does not set a domain for the `LtpaToken` cookie.

If you specify the `UseDomainFromURL` value, the application server sets the SSO domain name value to the domain of the host that is used in the Web address. For example, if an HTTP request comes from `server1.raleigh.ibm.com`, the application server sets the SSO domain name value to `raleigh.ibm.com`.

Note: The `UseDomainFromURL` value is case insensitive. You can type `usedomainfromurl` to use this value.

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

Interoperability mode:

Specifies that an interoperable cookie is sent to the browser to support back-level servers.

In WebSphere Application Server, Version 6 and later, a new cookie format is needed by the security attribute propagation functionality. When the interoperability mode flag is enabled, the server can send a maximum of two single sign-on (SSO) cookies back to the browser. In some cases, the server just sends the interoperable SSO cookie.

Web inbound security attribute propagation:

When Web inbound security attribution propagation is enabled, security attributes are propagated to front-end application servers. When this option is disabled, the single sign-on (SSO) token is used to log in and recreate the Subject from the user registry.

If the application server is a member of a cluster and the cluster is configured with a data replication service (DRS) domain, then propagation occurs. If DRS is not configured, then the SSO token contains the originating server information.

With this information, the receiving server can contact the originating server using an MBean call to get the original serialized security attributes.

com.tivoli.pd.jcfg.PDJrteCfg utility for Tivoli Access Manager single sign-on

The com.tivoli.pd.jcfg.PDJrteCfg utility configures the Java Runtime Environment component for Tivoli Access Manager. This component enables WebSphere Application Server to use Tivoli Access Manager security.

Purpose

Note: If you are using the WebSphere Application Server Network Deployment product, run the pdjrtecfg utility first on the deployment manager. Then, run the script on the other nodes in the cell.

Syntax

```
java com.tivoli.pd.jcfg.PDJrteCfg -action {config | unconfig} -host policy_server_host -was -java_home jre_path
```

Parameters

-action {config|unconfig}

Specifies the action to be performed. Actions include:

config Use to configure the Access Manager Java Runtime Environment component.

unconfig

Use to reconfigure the Access Manager Java Runtime Environment component.

-cfgfiles_path

Specifies where the generated configuration files will be placed.

Note: This parameter is required.

-host *policy_server_host*

Specifies the policy server host name.

Valid values for *policy_server_host* include any valid IP host name.

Examples include:

```
host = libra
```

```
host = libra.dallas.ibm.com
```

-was

Notifies Tivoli Access Manager Runtime for Java that the WebSphere Application Server version is

being configured so it is not necessary to perform certain steps such as copying the Java security jar files and PD.jar file since they were already placed in the appropriate directory by the WebSphere Application Server installer.

-java_home *jre_path*

Specifies the fully qualified path to the Java runtime (such as the directory ending in jre). If this parameter is not specified, the home directory for the jre in the PATH statement is used. If the home directory for the jre is not in the PATH statement, this utility can create an incorrect parameter in the output files.

Comments

This command copies Tivoli Access Manager Java libraries to a library extensions directory that exists for a Java runtime that has already been installed on the system.

You can install more than one Java Runtime Environment (JRE) on a given machine. The pdjrtecfg command can be used to configure the Tivoli Access Manager Java Runtime Environment component independently for each of the JRE configurations.

```

${JAVA_HOME}/bin/java
-Dfile.encoding=ISO8859-1 \
-Dws.output.encoding=CP1047 \
-Xnoargsconversion \
-Dpd.home=${WAS_HOME}/java/jre/PolicyDirector \
-cp ${WAS_HOME}/java/jre/lib/ext/PD.jar \
  com.tivoli.pd.jcfg.PDJrteCfg \
-action config \
  -cfgfiles_path ${WAS_HOME}/java/jre \
  -host gary.us.ibm.com \
  -java_home ${JAVA_HOME}
```

where:

-Dws.output.encoding

Is used to enable z/OS to display all of its messages and errors in a readable format.

-Dpd.home

Indicates where Tivoli Access Manager Runtime for Java has been installed. For WebSphere Application Server, this is java.home/PolicyDirector

Related information

Application Programming Interface documentation for IBM Tivoli Access Manager V2R1

The pdjrtecfg utility for IBM Tivoli Access Manager V2R1

com.tivoli.pd.jcfg.SvrSslCfg utility for Tivoli Access Manager single sign-on

The utility is used to configure and remove the configuration information associated with WebSphere Application Server and the Tivoli Access Manager server.

Purpose

Run the svrsslcfg script first on the deployment manager and then on the other nodes in the cell.

Syntax

```

java com.tivoli.pd.jcfg.SvrSslCfg
-action {config | unconfig} -admin_id admin_user_ID
-admin_pwd admin_password -appsvr_id application_server_name
-appsvr_pwd application_server_password -mode{local|remote}
-host host_name_of_application_server
```

```
-policysvr policy_server_name:port:rank [,...]
-authzsvr authorization_server_name:port:rank [,...]
-cfg_file fully_qualified_name_of_configuration_file
-domain Tivoli_Access_Manager_domain
-key_file fully_qualified_name_of_keystore_file
-cfg_action {create|replace}
```

Parameters

-action {config | unconfig}

Specifies the configuration action that is performed by the script. The following options apply:

-action config

Configuring a server creates user and server information in the user registry and creates local configuration and key store files on the application server. Use the `-action unconfig` option to reverse this operation.

If this action is specified, the following options are required: `-admin_id`, `-admin_pwd`, `-appsvr_id`, `-port`, `-mode`, `-policysvr`, `-authzsvr`, and `-key_file`.

-action unconfig

Reconfigures an application server to complete the following actions:

- Remove the user and server information from the user registry
- Delete the local key store file
- Remove information for this application from the configuration file without deleting the file

The reconfiguration operation fails only if the caller is unauthorized or the policy server cannot be contacted.

This action can succeed when a configuration file does not exist. When the configuration file does not exist, it is created and used as a temporary file to hold configuration information during the operation, and then the file is deleted completely.

If this action is specified, the following options are required: `-admin_id`, `-admin_pwd`, `-appsvr_id`, and `-policysvr`.

-admin_id *admin_user_ID*

Specifies the Tivoli Access Manager administrator name. If this option is not specified, `sec_master` is the default.

A valid administrative ID is an alphanumeric, case-sensitive string. String values are expected to be characters that are part of the local code set. You cannot use a space in the administrative ID.

For example, for U.S. English the valid characters are the letters a-Z, the numbers 0-9, a period (.), an underscore (_), a plus sign (+), a hyphen (-), an at sign (@), an ampersand (&), and an asterisk (*). The minimum and maximum lengths of the administrative ID, if there are limits, are imposed by the underlying registry.

-admin_password *admin_password*

Specifies the password of the Tivoli Access Manager administrator user that is associated with the `-admin_id` parameter. The password restrictions depend upon the password policy for your Tivoli Access Manager configuration.

-appsvr_id *application_server_name*

Specifies the name of the application server. The name is combined with the host name to create unique names for Tivoli Access Manager objects created for your application. The following names are reserved for Tivoli Access Manager applications: `ivacl`, `secmgrd`, `ivnet`, and `ivweb`.

-appsvr_pwd *application_server_password*

Specifies the password of the application server. This option is required. A password is created by the system and the configuration file is updated with the password created by the system.

If this option is not specified, the server password will be read from standard input.

-authzsvr *authorization_server_name*

Specifies the name of the Tivoli Access Manager authorization server with which the application server communicates. The server is specified by fully qualified host name, the SSL port number, and the rank. The default SSL port number is 7136. For example: myauth.mycompany.com:7136:1. You can specify multiple servers if the entries are separated by a comma (,).

-cfg_action {create | replace}

Specifies the action to take when creating the configuration and key files. Valid values are **create** or **replace**. Use the **create** option to initially create the configuration and keystore files. Use the **replace** option if these files already exist. If you use the **create** option and the configuration or keystore files already exist, an exception is created.

Options are as follows:

create Specifies to create the configuration and key store files during server configuration. Configuration fails if either of these files already exists.

replace

Specifies to replace the configuration and key store files during server configuration. Configuration deletes any existing files and replaces them with new ones.

-cfg_file *fully_qualified_name_of_configuration_file*

Specifies the configuration file path and name.

A file name should be an absolute file name (fully qualified file name) to be valid.

-domain *Tivoli_Access_Manager_domain*

Specifies the Tivoli Access Manager domain name to which the administrator is authenticated. This domain must exist and an the administrator ID and password must be valid for this domain. The application server is specified in this domain.

If not specified, the local domain that was specified during Tivoli Access Manager runtime configuration will be used. The local domain value will be retrieved from the configuration file.

A valid domain name is an alphanumeric, case-sensitive string. String values are expected to be characters that are part of the local code set. You cannot use a space in the domain name.

For example, for U.S. English the valid characters for domain names are the letters a-Z, the numbers 0-9, a period (.), an underscore (_), a plus sign (+), a hyphen (-), an at sign (@), an ampersand (&), and an asterisk (*). The minimum and maximum lengths of the domain name, if there are limits, are imposed by the underlying registry.

-host *host_name_of_application_server*

Specifies the TCP host name used by the Tivoli Access Manager policy server to contact this server. This name is saved in the configuration file using the `azn-app-host` key.

The default is the local host name returned by the operating system. Valid values for `host_name` include any valid IP host name.

Examples:

```
host = libra
host = libra.dallas.ibm.com
```

-key_file *fully_qualified_name_of_keystore_file*

Specifies the directory that is to contain the key files for the server. A valid directory name is determined by the operating system. Use a fully qualified file name that contains the application server certificate and key file.

Make sure that server user (for example, `ivmgr`) or all users have permission to access the `.kdb` file and the folder that contains the `.kdb` file.

This option is required.

-mode *server_mode*

Specifies the mode in which the application operates. This value must be either local or remote.

-policysvr *policy_server_name*

Specifies the name of the policy server.

Comments

After the successful configuration of a Tivoli Access Manager Java application server, SvrSslCfg creates a user account and server entries representing the Java application server in the Tivoli Access Manager user registry. In addition, SvrSslCfg creates a configuration file and a Java key store file, which securely stores a client certificate, locally on the application server. This client certificate permits callers to make authenticated use of Tivoli Access Manager services. Conversely, reconfiguration removes the user and server entries from the user registry and cleans up the local configuration and keystore files.

The contents of an existing configuration file can be modified by using the SvrSslCfg utility. The configuration file and the key store file must already exist when calling SvrSslCfg with all options other than -action config or -action unconfig.

The following options are parsed and processed into the configuration file, but are otherwise ignored in this version of Tivoli Access Manager:

The host name is used to build a unique name (identity) for the application. The pdadmin user list command displays the application identity name in the following format:

server_name/host_name

Note that the pdadmin server list command displays the server name in a slightly different format:

server_name-host_name

```
CLASSPATH=${WAS_HOME}/java/jre/lib/ext/PD.jar:${WAS_CLASSPATH}
java \
-cp ${CLASSPATH} \
-Dpd.cfg.home= ${WAS_HOME}/java/jre \
-Dfile.encoding=ISO8859-1 \
-Dws.output.encoding=CP1047 \
-Xnoargsconversion \
  com.tivoli.pd.jcfg.SvrSslCfg \
-action config \
-admin_id sec_master \
-admin_pwd $TAM_PASSWORD \
-appsvr_id $APPSVR_ID \
-policysvr ${TAM_HOST}:7135:1 \
-port 7135 \
-authzsvr ${TAM_HOST}:7136:1 \
-mode remote \
-cfg_file ${CFG_FILE} \
-key_file ${KEY_FILE} \
-cfg_action create
```

Creating a trusted user account in Tivoli Access Manager

Tivoli Access Manager trust association interceptors require the creation of a trusted user account in the shared LDAP user registry.

About this task

This account includes the ID and password that WebSEAL uses to identify itself to WebSphere Application Server. To prevent potential vulnerabilities, do not use the `sec_master` ID as the trusted user account and ensure that the password you use is unique and generated randomly. Use the trusted user account for the TAI or TAI++ only.

1. Use either the Tivoli Access Manager `pdadmin` command-line utility or Web Portal Manager to create the trusted user. For example, from the **pdadmin** command line.
2. Reference the code listed below as an example for creating a trusted user account.
3. Reference the following additional resources for more information:
 - a. “Configuring WebSEAL for use with WebSphere Application Server”
 - b. “Configuring Tivoli Access Manager plug-in for Web servers for use with WebSphere Application Server” on page 362

Example

```
pdadmin> user create webseal_userid webseal_userid_DN firstname  
surname password
```

```
pdadmin> user modify webseal_userid account-valid yes
```

Configuring WebSEAL for use with WebSphere Application Server

Use this topic to set the SSO password in WebSEAL for single sign-on to WebSphere Application Server.

About this task

A junction must be created between WebSEAL and WebSphere Application Server. This junction carries the `iv-credentials` (for TAI++) or `iv-user` (for TAI) and the HTTP basic authentication headers with the request. You can configure WebSEAL to pass the end user identity in other ways, the `iv-credentials` header is the only one supported by the TAI++ and the `iv-user` is the only one supported by TAI.

We recommend that communications over the junction use Secure Sockets Layer (SSL) for increased security. Setting up SSL across this junction requires that you configure the HTTP Server used by WebSphere Application Server, and WebSphere Application Server itself, to accept inbound SSL traffic and route it correctly to WebSphere Application Server. This activity requires importing the necessary signing certificates into the WebSEAL certificate keystore, and possibly also the HTTP Server certificate keystore.

Create the junction between WebSEAL and WebSphere Application Server using the **-c iv_creds** option for TAI++ and **-c iv_user** for TAI. Enter either of the following commands as one line using the variables that are appropriate for your environment:

TAI++

```
server task webseald-server create -t ssl -b supply -c iv_creds  
-h host_name -p websphere_app_port_number junction_name
```

TAI

```
server task webseald-server create -t ssl -b supply -c iv_user  
-h host_name -p websphere_app_port_number junction_name
```

Note:

1. If warning messages are displayed about the incorrect setup of certificates and key databases, delete the junction, correct problems with the key databases, and recreate the junction.
2. The junction can be created as `-t tcp` or `-t ssl`, depending on your requirements.

For single sign-on (SSO) to WebSphere Application Server the SSO password must be set in WebSEAL. To set the password, complete the following steps:

1. Edit the WebSEAL configuration file `webseal_install_directory/etc/webseald-default.conf` Set the following parameter: `basicauth-dummy-passwd=webseal_userid_passwd` where `webseal_userid_passwd` is the SSO password for the trusted user account set in “Creating a trusted user account in Tivoli Access Manager” on page 360.
2. Restart WebSEAL.

What to do next

For more details and options about how to configure junctions between WebSEAL and WebSphere Application Server, including other options for specifying the WebSEAL server identity, refer to the *Tivoli Access Manager WebSEAL Administration Guide* as well as to the documentation for the HTTP Server you are using with your WebSphere Application Server. Tivoli Access Manager documentation is available at <http://publib.boulder.ibm.com/tividd/td/tdprodlst.html>.

Configuring Tivoli Access Manager plug-in for Web servers for use with WebSphere Application Server

Tivoli Access Manager plug-in for Web servers can be used as a security gateway for your protected WebSphere Application Server resources.

About this task

With such an arrangement the plug-in authorizes all user requests before passing the credentials of the authorized user to WebSphere Application Server in the form of an iv-creds header. Trust between the plug-in and WebSphere Application Server is established through use of basic authentication headers containing the single sign-on (SSO) user password.

1. The Tivoli Access Manager plug-in for Web servers configuration shows IV headers configured for post-authorization processing, and basic authentication that is configured as the authentication mechanism and for post-authorization processing, as shown in the example below.
2. After a request is authorized, the basic authentication header is removed from the request (`strip-hdr=always`) and a new one is added (`add-hdr=supply`).
3. Included in this new header is the password that is set when the SSO user is created in “Creating a trusted user account in Tivoli Access Manager” on page 360.
4. Specify this password in the **supply-password** parameter and it is passed in the newly created header. This basic authentication header enables trust between WebSphere Application Server and the plug-in.
5. An iv-creds header is also added (`generate=iv-creds`), which contains the credential information of the user passed onto WebSphere Application Server. Session cookies are used to maintain session state.

Example

```
[common-modules]
authentication = BA
session = session-cookie
post-authzn = BA
post-authzn = iv-headers
```

```
[iv-headers]
accept = all
generate = iv-creds
```

```
[BA]
strip-hdr = always
add-hdr = supply
supply-password = sso_user_password
```

What to do next

“Configuring single sign-on using trust association” or “Configuring single sign-on using trust association interceptor ++”

Configuring single sign-on using trust association

This task is performed to enable single sign-on using trust association. Trust association is used to connect reversed proxy servers to the application server.

Before you begin

Note: Use of TAIs for Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) authentication is deprecated in this release. The SPNEGO Web authentication panels provide a much easier and less error-prone way to configure SPNEGO.

To establish the trust association for the single sign-on, perform the following steps:

1. From the administrative console for WebSphere Application Server, click **Security > Global security**.
2. From Authentication mechanisms, click **Web and SIP security > Trust association**.
3. Select the **Enable trust association** option.
4. Under Additional properties, click the **Interceptors** link.
5. Click **com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus** to use a WebSEAL interceptor, or **com.ibm.ws.security.spnego.TrustAssociationInterceptorImpl** to use a SPNEGO interceptor.
6. Under Custom properties, select a custom property to edit or click **New** to create a new one. Enter the property name and value pairs.
7. Click **OK**.
8. Save the configuration and log out.
9. Restart WebSphere Application Server.

Configuring single sign-on using trust association interceptor ++

Perform this task to enable single sign-on using trust association interceptor ++. The steps involve setting up trust association and creating the interceptor properties.

Before you begin

Although you can use Simple WebSphere Authentication Mechanism (SWAM) by selecting the **Use SWAM-no authenticated communication between servers** option on the Authentication mechanisms and expiration panel, single sign-on (SSO) requires LTPA as the configured authentication mechanism.

To establish the trust association for the single sign-on, perform the following steps:

1. Click **Enable Trust Association**.
2. Click **Interceptors**.
3. Click **com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus** to use a WebSEAL interceptor. This interceptor is one of two WebSEAL interceptors that are supplied for your use. You choose to use this interceptor by supplying properties as described in the next step.

Note: WebSphere Application Server attempts to initialize both of these interceptors even if you only supplied properties for the **com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus** interceptor. As a result, messages AWXRB0008E and SECJ0384E can appear during initialization to indicate that the interceptor you did not choose has failed to initialize. This is normal processing and does not affect the initialization of the interceptor you did select. To

inhibit the display of messages AWXRB0008E and SECJ0384E, you can delete the interceptor you do not want to use prior to beginning the initialization. You can add that interceptor back later if your environment changes.

4. Click **Custom Properties**.
5. Click **New** to enter the property name and value pairs. Ensure that the following parameters are set:

Table 15. Custom properties

Option	Description
com.ibm.websphere.security.webseal.checkViaHeader	<p>You can configure TAI so that the via header can be ignored when validating trust for a request. Set this property to <i>false</i> if none of the hosts in the via header need to be trusted. When set to <i>false</i> you do not need to set the trusted host names and host ports properties. The only mandatory property to check when via header is <i>false</i> is com.ibm.websphere.security.webseal.loginId.</p> <p>The default value of the check via header property is <i>false</i>. When using Tivoli Access Manager plug-in for Web servers, set this property to <i>false</i>.</p> <p>Note: The via header is part of the standard HTTP header that records the server names the request that passed through.</p>
com.ibm.websphere.security.webseal.loginId	<p>The WebSEAL trusted user as created in “Creating a trusted user account in Tivoli Access Manager” on page 360 The format of the username is the short name representation. This property is mandatory. If it is not set in WebSphere Application Server, the TAI initialization fails.</p>
com.ibm.websphere.security.webseal.id	<p>A comma-separated list of headers that exists in the request. If all of the configured headers do not exist in the request, trust cannot be established. The default value for the ID property is <i>iv-creds</i>. Any other values set in WebSphere Application Server are added to the list along with <i>iv-creds</i>, separated by commas.</p>
com.ibm.websphere.security.webseal.hostnames	<p>Do not set this property if using Tivoli Access Manager Plug-in for Web Servers. The property specifies the host names (case sensitive) that are trusted and expected in the request header. Requests arriving from un-listed hosts might not be trusted. If the checkViaHeader property is not set or is set to false then the trusted host names property has no influence. If the checkViaHeader property is set to true, and the trusted host names property is not set, TAI initialization fails.</p>
com.ibm.websphere.security.webseal.ports	<p>Do not set this property if using Tivoli Access Manager plug-in for Web servers. This property is a comma-separated list of trusted host ports. Requests that arrive from unlisted ports might not be trusted. If the checkViaHeader property is not set, or is set to false this property has no influence. If the checkViaHeader property is set to true, and the trusted host ports property is not set in WebSphere Application Server, the TAI initialization fails.</p>

Table 15. Custom properties (continued)

Option	Description
com.ibm.websphere.security.webseal.viaDepth	<p>A positive integer that specifies the number of source hosts in the via header to check for trust. By default, every host in the via header is checked, and if any host is not trusted, trust cannot be established. The via depth property is used when only some of the hosts in the via header have to be trusted. The setting indicates the number of hosts that are required to be trusted.</p> <p>As an example, consider the following header: Via: HTTP/1.1 webseal1:7002, 1.1 webseal2:7001</p> <p>If the viaDepth property is not set, is set to 2 or is set to 0, and a request with the previous via header is received then both webseal1:7002 and webseal2:7001 need to be trusted. The following configuration applies: com.ibm.websphere.security.webseal.hostnames = webseal1,webseal2 com.ibm.websphere.security.webseal.ports = 7002,7001</p> <p>If the via depth property is set to 1, and the previous request is received, then only the last host in the via header needs to be trusted. The following configuration applies: com.ibm.websphere.security.webseal.hostnames = webseal2 com.ibm.websphere.security.webseal.ports = 7001</p> <p>The viaDepth property is set to 0 by default, which means all of the hosts in the via header are checked for trust.</p>
com.ibm.websphere.security.webseal.ssoPwdExpiry	<p>After trust is established for a request, the single sign-on user password is cached, eliminating the need to have the TAI re-authenticate the single sign-on user with Tivoli Access Manager for every request. You can modify the cache timeout period by setting the single sign-on password expiry property to the required time in seconds. If the password expiry property is set to 0, the cached password never expires. The default value for the password expiry property is 600.</p>
com.ibm.websphere.security.webseal.ignoreProxy	<p>This property can be used to tell the TAI to ignore proxies as trusted hosts. If set to true the comments field of the hosts entry in the via header is checked to determine if a host is a proxy. Remember that not all proxies insert comments in the via header indicating that they are proxies. The default value of the ignoreProxy property is false. If the checkViaHeader property is set to false then the ignoreProxy property has no influence in establishing trust.</p>
com.ibm.websphere.security.webseal.configURL	<p>For the TAI to establish trust for a request, it requires that the SvrSslCfg run for the Java Virtual Machine on the Application Server and result in the creation of a properties file. If this properties file is not at the default URL, which is file://java.home/PdPerm.properties, the correct URL of the properties file must be set in the configuration URL property. If this property is not set, and the SvrSslCfg-generated properties file is not in the default location, the TAI initialization fails. The default value for the config URL property is file://\${WAS_INSTALL_ROOT}/java/jre/PdPerm.properties.</p>

6. Click **OK**.
7. Save the configuration and log out.
8. Restart WebSphere Application Server.

Configuring global sign-on principal mapping

You can create a new application login that uses the Tivoli Access Manager GSO database to store the login credentials.

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > Application logins**.

3. Click **New** to create a new Java Authentication and Authorization Service (JAAS) login configuration.
4. Enter the alias name of the new application login. Click **Apply**.
5. Under Additional properties, click **JAAS login modules** to define the JAAS Login Modules.
6. Click **New** and enter the following information:

Module class name: com.tivoli.pdwas.gso.AMPrincipalMapper
Use Login Module Proxy: enable
Authentication strategy: REQUIRED

7. Click **Apply**
8. Under Additional Properties section, click **Custom Properties** to define login module-specific values that are passed directly to the underlying login modules.
9. Click **New**.

The Tivoli Access Manager principal mapping module uses the authDataAlias configuration string to retrieve the correct user name and password from the security configuration.

The authDataAlias attribute that is passed to the module is configured for the J2C connection factory. Because the authDataAlias attribute is an arbitrary string that is entered at configuration time, the following scenarios are possible:

- The authDataAlias attribute contains both the global sign-on (GSO) resource name and the user name. The format of this string is "Resource/User".
- The authDataAlias attribute contains the GSO Resource name only. The user name is determined by using the Subject of the current session.

The scenario to use is determined by a JAAS configuration option, as shown here:

Name: com.tivoli.pd.as.gso.AliasContainsUserName

Value: True, if the alias contains the user name; false, if the user name must be retrieved from the security context

When entering authDataAlias attributes through the WebSphere Application Server administrative console, the node name is automatically pre-pended to the alias. The JAAS configuration entry determines whether this node name is removed or included as part of the resource name, as shown here:

Name: com.tivoli.pd.as.gso.AliasContainsNodeName

Value: True, if the alias contains the node name

Note: If the PdPerm.properties configuration file is not located in the JAVA_HOME/PdPerm.properties default location, then you also need to add the following property:

Name: com.tivoli.pd.as.gso.AMCfgURL

Value: file:///path to PdPerm.properties

Enter each new parameter using the following scenario information as a guide, then click **Apply**.

Scenario 1

Auth Data Alias - BackendEIS/eisUser

Resource - BackEndEIS

User - eisUser

Principal Mapping Parameters

Name	Value
delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	true
com.tivoli.pd.as.gso.AliasContainsNodeName	false
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path

debug	false
-------	-------

Scenario 2

Auth Data Alias - BackendEIS

Resource - BackEndEIS

User - Currently authenticated WebSphere Application Server user

Principal Mapping Parameters

Name	Value
delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	false
com.tivoli.pd.as.gso.AliasContainsNodeName	false
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path
debug	false

Scenario 3

Auth Data Alias - nodename/BackendEIS/eisUser

Resource - BackEndEIS

User - eisUser

Principal Mapping Parameters

Name	Value
delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	true
com.tivoli.pd.as.gso.AliasContainsNodeName	true
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path
debug	false

Scenario 4

Auth Data Alias - nodename/BackendEIS/eisUser

Resource - nodename/BackEndEIS (notice that node name is not removed)

User - eisUser

Principal Mapping Parameters

Name	Value
delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	true
com.tivoli.pd.as.gso.AliasContainsNodeName	false
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path
debug	false

Scenario 5

Auth Data Alias - BackendEIS/eisUser

Resource - BackEndEIS

User - eisUser

Principal Mapping Parameters

Name	Value
------	-------

delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	false
com.tivoli.pd.as.gso.AliasContainsNodeName	true
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path
debug	false

Scenario 6

Auth Data Alias - nodename/BackendEIS/eisUser

Resource - nodename/BackendEIS/eisUser

(notice that the resource is the same as Auth Data Alias).

User - Currently authenticated WebSphere Application Server user

Principal Mapping Parameters

Name	Value
delegate	com.tivoli.pdwas.gso.AMPrincipalMapper
com.tivoli.pd.as.gso.AliasContainsUserName	false
com.tivoli.pd.as.gso.AliasContainsNodeName	false
com.tivoli.pd.as.gso.AMLoggingURL	file:///jlog_props_path
debug	false

10. Create the Java 2 Connector (J2C) authentication aliases. The user name and password that are assigned to these alias entries are irrelevant because Tivoli Access Manager is responsible for providing user names and passwords. However, the user name and password that are assigned to the J2C authentication aliases need to exist so that they can be selected for the J2C connection factory in the administrative console.

To create the J2C authentication aliases, from the WebSphere Application Server administrative console, click **Security > Global security**. Under Authentication, click **Java Authentication and Authorization Service > J2C authentication data**, and then click **New** for each new entry. Refer to the previous table for scenario inputs.

The connection factories for each resource adapter that need to use the GSO database must be configured to use the Tivoli Access Manager Principal mapping module:

- From the WebSphere Application Server administrative console, click **Applications > Enterprise Applications > application_name > Resource references**. Note that J2C connection factories must be already configured for the selected application. To configure a new J2C connection factory, see tdat_confconfac.dita.
- Under Additional properties, click **Resource Adapter**.
The resource adapter can be standalone and does not need to be packaged with the application. The resource adapter is configured from **Resources > Resource Adapters** for standalone scenarios.
- Under Additional properties, click **J2C Connection Factories**.
- Click **New** and enter the connection factory properties.
- When finished, click **Apply > Save**.

Note:

Custom mapping configuration for the connection factory is deprecated in WebSphere Application Server Version 6. To configure the GSO credential mapping, use the Map Resource References to Resources panel on the administrative console. For more information, see J2EE connector security.

Configuring administrative authentication

An authentication mechanism defines rules about security information, such as whether a credential is forwardable to another Java process, and the format of how security information is stored in both credentials and tokens. The Rivest Shamir Adleman (RSA) token authentication mechanism simplifies the security environment for flexible management topology, that is, the topology where you can locally or remotely submit and manage administrative jobs through a job manager that manages applications, perform product maintenance, modify configurations, and control the application server runtime. You use the administrative console to configure administrative authentication, which involves the configuring of the Rivest Shamir Adleman (RSA) token authentication mechanism.

Before you begin

The following keystore, truststore, and rootstore descriptions give you an idea of where certificates are stored and how trust is configured between processes.

The **NodeRSATokenKeyStore** contains the Rivest Shamir Adleman (RSA) token personal certificate used for this process. Not only is the public/private key from this certificate used to create RSA tokens, but the public key is used by other processes to create tokens. The RSA personal certificate is signed by an RSA root certificate.

The **NodeRSATokenTrustStore** contains all RSA signer certificates from other processes that are trusted to send RSA tokens to this process. The signers in this trust store are placed there automatically during the registration process. However, this task allows an administrator to configure trust between to processes not normally involved in the same administrative domain. There may be requirements where two base servers are communicating administratively. When using the RSA token authentication mechanism, the base servers need to share RSA signers if administrative communications is operating in both directions.

The **NodeRSATokenRootStore** contains the root personal certificate that is used to create new RSA personal certificates. Do not use the root certificate to create RSA tokens because this usage compromises the long-lived keys. Only use the root certificate to sign other certificates.

No manual steps are required with these keystores, and this allows uncommon trust establishment among processes not in the same administrative domain. You can also replace the RSA personal certificate with a personal certificate obtained from a certificate authority (CA) if desired. In this case, make sure the CA root certificate is placed in all RSA trust stores in the same administrative domain.

1. Click **SSL certificate and key management > Global security** . Under Keystore usages select **RSA token keystores** from the drop-down list.
2. Select the RSA token key store you want to administer.
3. Modify the description if required.
4. Modify the path if required.
5. Select **read only, initialize at setup**, or both if required.
6. Enter the correct password to make these modifications
7. Click **Apply** and **Save**.

Results

You configured administrative authentication.

What to do next

In cases where the process is back-level or a target RSA certificate cannot be obtained, the fallback mechanism is Lightweight Third-Party Authentication (LTPA) which is supported in all previous releases for

administrative communications. The fallback occurs automatically. If the LTPA keys are not shared and a fallback occurs, LTPA will fail as well. However, this situation is typically an error case in the RSA mechanism and should occur infrequently.

Job manager security

When performing a job manager registration process there are a number of WebSphere Application Server security impacts to consider.

Note: In this release of WebSphere Application Server, a new style of system management called *flexible management* has been introduced. It differs from the existing style of synchronous invocation and response calls through wsadmin or Java APIs by offering an asynchronous job queuing mechanism for administration purposes. At the core of flexible management is a new administrative process called the *job manager*. You can make both application servers registered to administrative agents and deployment manager servers known to the job manager through a registration process. After you register the servers, you can queue administrative jobs directed at the application servers or deployment managers through the job manager. You can submit these jobs to a large number of servers over a geographically dispersed area. There are a number of security considerations you must keep in mind both during and after the job manager registration process.

Consider the following:

- Security configuration requirements should be kept to a minimum prior to registration.
- Allow an agent or a deployment manager (dmgr) to be federated to the job manager with almost any security configuration. Some exceptions include:
 - The administrative agent or deployment manager must have the same administrative security state (either enabled or disabled).
 - To enable security after federation, you must enable all administrative agent and deployment manager processes within the same administrative domain, then restart all of the processes at the same time.
- Leverage the creation of a chained certificate to exchange only the long-lived root certificates between an administrative agent, deployment manager and job manager. When a personal certificate expires in either the administrative agent, deployment manager or the job manager, it does not affect trust that was established during federation.
- Use the Rivest Shamir Adleman (RSA) certificate administrative-specific authentication mechanism, which does not rely on shared keys and is the default administrative authentication mechanism for the job manager. The RSA token authentication mechanism is also new to this release of WebSphere Application Server. Read about “RSA token authentication mechanism” on page 282 for more information.
- Add a profile Universal Unique Identifier (UUID) to all certificates generated in WebSphere Application Server Version 7.0. This profile UUID is used to authorize requests to extract jobs from the job manager queue.
- Job Manager jobs can be associated with caller credentials; either Lightweight Third-Party Authentication (LTPA) or Kerberos, or with specified credentials (the userid and password). Both are stored with the job. The password is obfuscated using the standard utilities and can be encrypted when the password encryption plug point is enabled. LTPA and Kerberos are refreshed as long as the authentication mechanism allows them to be.
- Administrative agent or deployment manager access to FileTransferServlet is performed by sending a valid RSA certificate that is trusted by the job manager and is validated by CertPath.

The required administrative roles for executing flexible management jobs are defined by the underlying administrative commands used by those jobs. For example, the required role for starting and stopping servers is the operator role. The operator role is also required for execution of the flexible management jobs that start and stop servers. The general rules for assigning required administrative roles are:

- Viewing data requires the monitor role.
- Updating data requires the configurator role.
- Managing jobs requires the operator role.
- Registering or un-registering managed nodes requires the administrator role.

Java Authentication and Authorization Service

The standard Java 2 security application programming interface (API) helps enforce access control based on the location of the code source or the author or packager of the code that signed the jar file. The current principal of the running thread is not considered in the Java 2 security authorization. Instances where authorization is based on the principal, as opposed to the code base, and the user exist. The Java Authentication and Authorization Service is a standard Java API that supports the Java 2 security authorization to extend the code base on the principal as well as the code base and users.

The Java Authentication and Authorization Service (JAAS) Version 1.0 extends the Java 2 security architecture of the Java 2 platform with additional support to authenticate and enforce access control with principals and users. JAAS implements a Java version of the standard Pluggable Authentication Module (PAM) framework, and extends the access control architecture of the Java 2 platform in a compatible fashion to support user-based authorization or principal-based authorization. WebSphere Application Server fully supports the JAAS architecture. JAAS extends the access control architecture to support role-based authorization for Java Platform, Enterprise Edition (Java EE) resources including servlets, JavaServer Pages (JSP) files, and Enterprise JavaBeans (EJB) components.

Refer to “Java 2 security” on page 44 for more information.

The following sections cover the JAAS implementation and programming model:

- “Login configuration for Java Authentication and Authorization Service” on page 406
- “Programmatic login” on page 393
- “Java Authentication and Authorization Service authorization”

The JAAS documentation can be found at <http://www.ibm.com/developerworks/java/jdk/security>. Scroll down to find the JAAS documentation for your platform.

Java Authentication and Authorization Service authorization

Java 2 security architecture uses a security policy to specify which access rights are granted to running code. This architecture is *code-centric*. The permissions are granted based on code characteristics including where the code is coming from, whether it is digitally signed, and by whom. Authorization of the Java Authentication and Authorization Service (JAAS) augments the existing code-centric access controls with new user-centric access controls. Permissions are granted based on what code is running and who is running it.

When using JAAS authentication to authenticate a user, a *subject* is created to represent the authenticated user. A subject is comprised of a set of principals, where each principal represents an identity for that user. You can grant permissions in the policy to specific principals. After the user is authenticated, the application can associate the subject with the current access control context. For each subsequent security-checked operation, the Java runtime automatically determines whether the policy grants the required permission to a specific principal only. If so, the operation is supported if the subject that is associated with the access control context contains the designated principal only.

Associate a subject with the current access control context by calling the static `doAs` method from the subject class, passing it an authenticated subject and the `java.security.PrivilegedAction` or `java.security.PrivilegedExceptionAction` method. The `doAs` method associates the provided subject with the current access control context and then invokes the `run` method from the action. The `run` method implementation contains all the code that ran as the specified subject. The action runs as the specified subject.

In the Java 2 Platform, Enterprise Edition (J2EE) programming model, when invoking the Enterprise JavaBeans (EJB) method from an enterprise bean or servlet, the method runs under the user identity that is determined by the run-as setting. The J2EE Version 1.4 Specification does not indicate which user identity to use when invoking an enterprise bean from a Subject.doAs action block within either the EJB code or the servlet code. A logical extension is to use the proper identity that is specified in the subject when invoking the EJB method within the Subject doAs action block.

Letting the Subject.doAs action overwrite the run-as identity setting is an ideal way to integrate the JAAS programming model with the J2EE run-time environment. However, JAAS introduced an issue into the Software Development Kit (SDK), Java Technology Edition Versions 1.3 or later when integrating the JAAS Version 1.0 or later implementation with the Java 2 security architecture. A subject, which is associated with the access control context is cut off by a doPrivileged call when a doPrivileged call occurs within the Subject.doAs action block. Until this problem is corrected, no reliable and run-time efficient way is available to guarantee the correct behavior of Subject.doAs action in a J2EE run-time environment.

The problem can be explained better with the following example:

```
Subject.doAs(subject, new java.security.PrivilegedAction() {
    Public Object run() {
        // Subject is associated with the current thread context
        java.security.AccessController.doPrivileged( new
            java.security.PrivilegedAction() {
                public Object run() {
                    // Subject was cut off from the current
                    // thread context

                }
            }
        );
        return null;
    }
});
// Subject is associated with the current thread context
return null;
}
```

In the previous code example, the Subject object is associated with the context of the current thread. Within the run method of a doPrivileged action block, the Subject object is removed from the thread context. After leaving the doPrivileged block, the Subject object is restored to the current thread context. Because doPrivileged blocks can be placed anywhere along the running path and instrumented quite often in a server environment, the run-time behavior of a doAs action block becomes difficult to manage.

To resolve this difficulty, WebSphere Application Server provides a WSSubject helper class to extend the JAAS authorization to a J2EE EJB method invocation, as described previously. The WSSubject class provides static doAs and doAsPrivileged methods that have identical signatures to the subject class. The WSSubject.doAs method associates the Subject to the currently running thread. The WSSubject.doAs and WSSubject.doAsPrivileged methods then invoke the corresponding Subject.doAs and Subject.doAsPrivileged methods. The original credential is restored and associated with the running thread upon leaving the WSSubject.doAs and WSSubject.doAsPrivileged methods.

The WSSubject class is not a replacement of the subject object, but rather a helper class to ensure consistent run-time behavior as long as an EJB method invocation is a concern.

The following example illustrates the run-time behavior of the WSSubject.doAs method:

```
WSSubject.doAs(subject, new java.security.PrivilegedAction() {
    Public Object run() {
        // Subject is associated with the current thread context
        java.security.AccessController.doPrivileged( new
```

```

        java.security.PrivilegedAction() {
            public Object run() {
                // Subject was cut off from the current thread
                // context.
            }
        }
    };
    // Subject is associated with the current thread context
    return null;
}
});

```

The `Subject.doAs` and `Subject.doAsPrivileged` methods are not integrated with the J2EE run-time environment. EJB methods that are invoked within the `Subject.doAs` and `Subject.doAsPrivileged` action blocks run under the identity that is specified by the run-as setting and not by the subject identity.

- The `Subject` object that is generated by the `WSLoginModuleImpl` instance and the `WSCClientLoginModuleImpl` instance contains a principal that implements the `WSPPrincipal` interface. Using the `getCredential` method for a `WSPPrincipal` object returns an object that implements the `WSCredential` interface. You can also find the `WSCredential` object instance in the `PublicCredentials` list of the subject instance. Retrieve the `WSCredential` object from the `PublicCredentials` list instead of using the `getCredential` method.
- The `getCallerPrincipal` method for the `WSSubject` class returns a string that represents the caller security identity. The return type differs from the `getCallerPrincipal` method of the `java.security.Principal` `EJBContext` interface.
- The `Subject` object that is generated by the Java 2 Connector (J2C) `DefaultPrincipalMapping` module contains a resource principal and a `PasswordCredentials` list. The resource principal represents the `RunAs` identity.

For more information, see J2EE connector security.

Related tasks

Chapter 5, “Authenticating users,” on page 111

The process of authenticating users involves a user registry and an authentication mechanism. Optionally, you can define trust between WebSphere Application Server and a proxy server, configure single sign-on capability, and specify how to propagate security attributes between application servers.

Using the Java Authentication and Authorization Service programming model for Web authentication

WebSphere Application Server supports the Java Platform, Enterprise Edition (Java EE) declarative security model. You can define the authentication and access control policy using the Java EE deployment descriptor. You can further stack custom login modules to customize the WebSphere Application Server authentication mechanism.

Before you begin

A custom login module can perform principal and credential mapping, custom security token and custom credential-processing, and error-handling among other possibilities. Typically, you do not need to use application code to perform authentication function. Use the programming techniques that are described in this section if you have to perform authentication function in application code. For example, if you have applications that programmed to the `SSOAuthenticator` helper function, you can use the following programming interface. The `SSOAuthenticator` helper function was deprecated starting with WebSphere Application Server Version 4.0. Use declarative security as a rule; use the techniques that are described in this section as a last resort.

About this task

When the Lightweight Third-Party Authentication (LTPA) mechanism single sign-on (SSO) option is enabled, the Web client login session is tracked by an LTPA SSO token cookie after successful login. At logout, this token is deleted to terminate the login session, but the server-side subject is not deleted. When you use the declarative security model, the WebSphere Application Server Web container performs client authentication and login session management automatically. You can perform authentication in application code by setting a login page without a Java EE security constraint and by directing client requests to your login page first. Your login page can use the Java Authentication and Authorization Service (JAAS) programming model to perform authentication. To enable WebSphere Application Server Web login modules to generate SSO cookies, use the following steps.

1. Create a new system login JAAS configuration. To access the panel, click **Security > Global security**. Under Java Authentication and Authorization Service, click **System logins**.
2. Manually clone the WEB_INBOUND login configuration, and give it a new alias. To clone the login configuration, click **New**, enter a name for the configuration, click **Apply**, then click **JAAS login modules** under Additional properties. Click **New** and configure the JAAS login module. For more information, see “Login module settings for Java Authentication and Authorization Service” on page 417. WebSphere Application Server Web container uses the WEB_INBOUND login configuration to authenticate Web clients. Changing the WEB_INBOUND login configuration affects all Web applications in the cell. You should create your own login configuration by cloning the contents of the WEB_INBOUND login configuration.
3. Select the `wsMapDefaultInboundLoginModule` login module and click **Custom properties**. There are two login modules defined in your login configuration: `ltpaLoginModule` and `wsMapDefaultInboundLoginModule`.
4. Add a login property name `cookie` with a value of **true**. The two login modules are enabled to generate LTPA SSO cookies. Do not add the cookie login option to the original WEB_INBOUND login configuration.
5. Stack your custom LoginModule(s) in the new login configuration (optional).
6. Use your login page for programmatic login by perform a JAAS `LoginContext.login` using your newly defined login configuration. After a successful login, either the `ltpaLoginModule` or the `wsMapDefaultInboundLoginModule` generates an LTPA SSO cookie upon a successful authentication. Exactly which LoginModule generates the SSO cookie depends on many factors, including system authentication configuration and runtime condition (which is beyond the scope of this section).
7. Call the modified `WSSubject.setRunAsSubject` method to add the subject to the authentication cache. The subject must be a WebSphere Application Server JAAS subject created by LoginModule. Adding the subject to the authentication cache recreates a subject from SSO token.
8. Use your programmatic logout page to revoke SSO cookies by invoking the `revokeSSOCookies` method from the `WSSecurityHelper` class. The term cookies is used because WebSphere Application Server Release 5.1.1 (and later) release supports a new LTPA SSO token with a different encryption algorithm, but can be configured to generate the original LTPA SSO token for backward compatibility. Note that the subject is still in the authentication cache and only the SSO cookies are revoked.

Example

Use the following code sample to perform authentication:

Suppose you wrote a `LoginServlet.java`:

```
import com.ibm.wsspi.security.auth.callback.WSCallbackHandlerFactory;
import com.ibm.websphere.security.auth.WSSubject;

public Object login(HttpServletRequest req, HttpServletResponse res)
    throws ServletException {

    PrintWriter out = null;
    try {
```



```

    out = res.getWriter();
    res.setContentType("text/html");
} catch (java.io.IOException e){
    // Error handling
}

Subject subject = null;
try {
    LoginContext lc = new LoginContext("system.Your_login_configuration",
WSCallbackHandlerFactory.getInstance().getCallbackHandler(
userid, null, password, req, res, null));
    lc.login();
    subject = lc.getSubject();
    WSSubject.setRunAsSubject(subject);
} catch(Exception e) {
    // catch all possible exceptions if you want or handle them separately
    out.println("Exception in LoginContext login + Exception = " +
e.getMessage());
    throw new ServletException(e.getMessage());
}

```

The following is sample code to revoke the SSO cookies upon a programming logout:

The LogoutServlet.java:

```

public void logout(HttpServletRequest req, HttpServletResponse res,
Object retCreds) throws ServletException {
    PrintWriter out =null;
    try {
        out = res.getWriter();
        res.setContentType("text/html");
    } catch (java.io.IOException e){
        // Error Handling
    }
    try {
        WSSecurityHelper.revokeSSOCookies(req, res);
    } catch(Exception e) {
        // catch all possible exceptions if you want or handle them separately
        out.println("JAASLogoutServlet: logout Exception = " + e.getMessage());
        throw new ServletException(e);
    }
}

```

What to do next

For more information on JAAS authentication, refer to Developing programmatic logins with the Java Authentication and Authorization Service. For more information on the AuthenLoginModule login module, refer to Example: Customizing a server-side Java Authentication and Authorization Service authentication and login configuration.

Developing custom login modules for a system login configuration

For WebSphere Application Server, multiple Java Authentication and Authorization Service (JAAS) plug-in points exist for configuring system logins. WebSphere Application Server uses system login configurations to authenticate incoming requests, outgoing requests, and internal server logins.

About this task

Application login configurations are called by Java Platform, Enterprise Edition (Java EE) applications for obtaining a Subject that is based on specific authentication information. This login configuration enables the application to associate the Subject with a specific protected remote action. The Subject is picked up

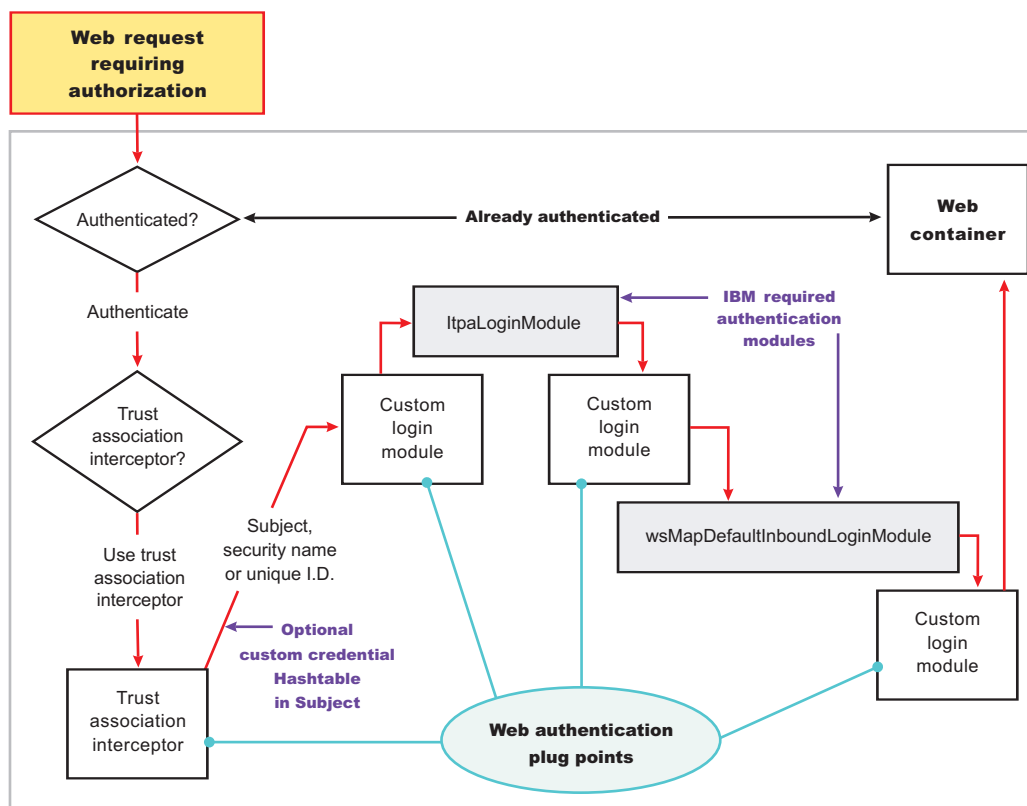
on the outbound request processing. The following list identifies the main system plug-in points. If you write a login module that adds information to the Subject of a system login, these are the main login configurations to plug in:

- WEB_INBOUND
- RMI_OUTBOUND
- RMI_INBOUND
- DEFAULT
- Authenticate Web requests with the WEB_INBOUND login configuration.

The WEB_INBOUND login configuration authenticates Web requests.

For more detailed information on the WEB_INBOUND configuration including its associated callbacks, see "RMI_INBOUND, WEB_INBOUND, DEFAULT" in "System login configuration entry settings for Java Authentication and Authorization Service" on page 409. Figure 1 shows an example of a configuration using a trust association interceptor (TAI) that creates a Subject with the initial information that is passed into the WEB_INBOUND login configuration. If the trust association interceptor is not configured, the authentication process goes directly to the WEB_INBOUND system login configuration, which consists of all the login modules combined in Figure 1. Figure 1 shows where you can plug in custom login modules and where the ItpaLoginModule and the wsMapDefaultInboundLoginModule login modules are required.

Figure 1



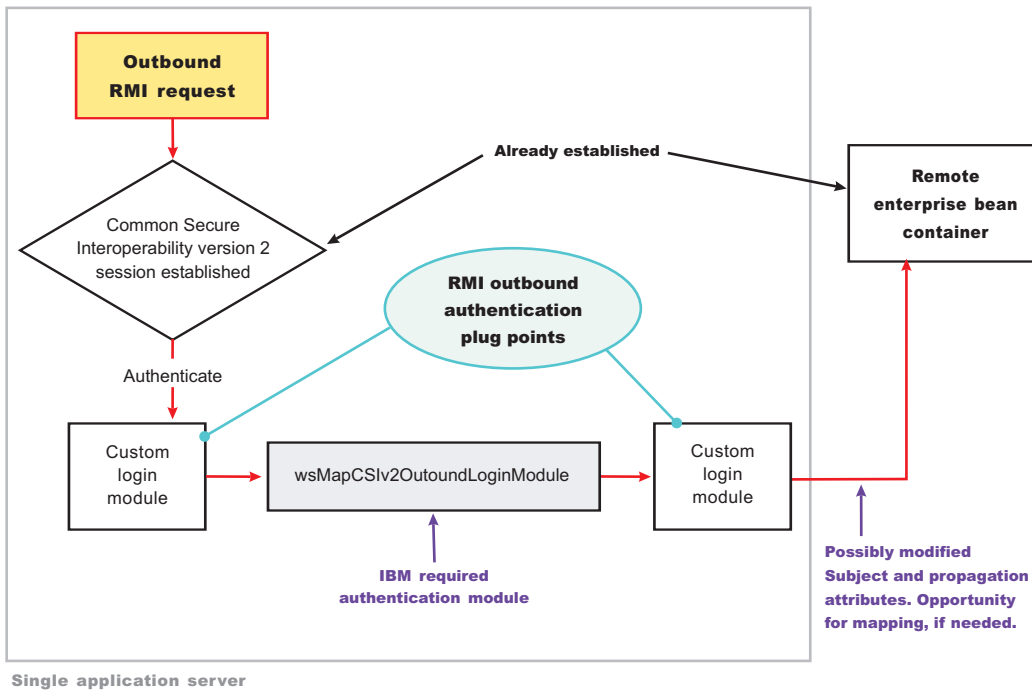
Single application server

- Handle outbound requests with the RMI_OUTBOUND login configuration.

The RMI_OUTBOUND login configuration is a plug point for handling outbound requests. WebSphere Application Server uses this plug point to create the serialized information that is sent downstream based on the invocation Subject passed in and other security context information such as propagation tokens. A custom login module can use this plug point to change the identity. For more information, see

“Configuring outbound mapping to a different target realm” on page 443. Figure 2 shows where you can plug in custom login modules and shows where the wsMapCSlv2OutboundLoginModule login module is required.


Figure 2



For more information on the RMI_OUTBOUND login configuration, including its associated callbacks, see “RMI_OUTBOUND” in “System login configuration entry settings for Java Authentication and Authorization Service” on page 409.

- Handle inbound authentication for enterprise bean requests with the RMI_INBOUND login configuration. The RMI_INBOUND login configuration is a plug point that handles inbound authentication for enterprise bean requests. WebSphere Application Server uses this plug point for either an initial login or a propagation login. For more information about these two login types, see “Security attribute propagation” on page 448. During a propagation login, this plug point is used to deserialize the information that is received from an upstream server. A custom login module can use this plug point to change the identity, handle custom tokens, add custom objects into the Subject, and so on. For more information on changing the identity using a Hashtable object, which is referenced in figure 3, see “Configuring inbound identity mapping” on page 435. Figure 3 shows where you can plug in custom login modules and shows that the ltpaLoginModule and the wsMapDefaultInboundLoginModule login modules are required.

Figure 3



For more information on the RMI_INBOUND login configuration, including its associated callbacks, see "RMI_INBOUND, WEB_INBOUND, DEFAULT" in "System login configuration entry settings for Java Authentication and Authorization Service" on page 409.

- Handle all other types of authentication requests with the DEFAULT login configuration. **DEFAULT login configuration**

The DEFAULT login configuration is a plug point that handles all of the other types of authentication requests, including administrative SOAP requests and internal authentication of the server ID. Propagation logins typically do not occur at this plug point.

For more information on the DEFAULT login configuration including its associated callbacks, see "RMI_INBOUND, WEB_INBOUND, DEFAULT" in "System login configuration entry settings for Java Authentication and Authorization Service" on page 409.

- Develop login configuration logic to know when specific information is present and how to use the information. **Writing a login module**

When you write a login module that plugs into a WebSphere Application Server application login or system login configuration, read the JAAS programming model, which is located at: <http://java.sun.com/products/jaas>. The JAAS programming model provides basic information about JAAS. However, before writing a login module for the WebSphere Application Server environment, read the following sections in this article:

- Useable callbacks
- Shared state variables
- Initial versus propagation logins
- Sample custom login module

Useable Callbacks

Each login configuration must document the callbacks that are recognized by the login configuration. However, the callbacks are not always passed data. The login configuration must contain logic to know when specific information is present and how to use the information. For example, if you write a custom login module that can plug into all four of the pre-configured system login configurations mentioned previously, three sets of callbacks might be presented to authenticate a request. Other callbacks might be present for other reasons, including propagation and making other information available to the login configuration.

Login information can be presented in the following combinations:

User name (NameCallback) and password (PasswordCallback)

This information is a typical authentication combination.

User name only (NameCallback)

This information is used for identity assertion, trust association interceptor (TAI) logins, and certificate logins.

Token (WSCredTokenCallbackImpl)

This information is for Lightweight Third Party Authentication (LTPA) token validation.

Propagation token list (WSTokenHolderCallback)

This information is used for a propagation login.

The first three combinations are used for typical authentication. However, when the WSTokenHolderCallback callback is present in addition to one of the first three information combinations, the login is called a *propagation login*. A propagation login means that some security attributes are propagated to this server from another server. The servers can reuse these security attributes if the authentication information validates successfully. In some cases, a WSTokenHolderCallback callback might not have sufficient attributes for a full login. Check the requiresLogin method on the WSTokenHolderCallback callback to determine if a new login is required. You can always ignore the information returned by the requiresLogin method, but, as a result, you might duplicate information. The following list contains the callbacks that might be present in the system login configurations. The list includes the callback name and a description of their responsibility.

callbacks[0] = new javax.security.auth.callback.NameCallback("Username: ");

This callback handler collects the user name for the login. The result can be the user name for a basic authentication login (user name and password) or a user name for an identity assertion login.

callbacks[1] = new javax.security.auth.callback.PasswordCallback("Password: ", false);

This callback handler collects the password for the login.

callbacks[2] = new

com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl("Credential Token: ");

This callback handler collects the Lightweight Third Party Authentication (LTPA) token or other token type for the login. This callback handler is typically present when a user name and password are not present.

callbacks[3] = new com.ibm.wsspi.security.auth.callback.WSTokenHolderCallback("Authz Token List: ");

This callback handler collects the ArrayList of TokenHolder objects that are returned from a call to the WSOpaqueTokenHelper.createTokenHolderListFromOpaqueToken API using the Common Secure Interoperability Version 2 (CSIv2) authorization token as input.

callbacks[4] = new

com.ibm.websphere.security.auth.callback.WSServletRequestCallback("HttpServletRequest: ");

This callback handler collects the HTTP servlet request object, if present. This callback handler enables login modules to get information from the HTTP request for use in the login, and is presented from the WEB_INBOUND login configuration only.

callbacks[5] = new

com.ibm.websphere.security.auth.callback.WSServletResponseCallback("HttpServletResponse: ");

This callback handler collects the HTTP servlet response object, if present. This callback handler enables login modules to put information into the HTTP response as a result of the login. An example of this situation might be adding the SingleSignonCookie cookie to the response. This callback handler is presented from the WEB_INBOUND login configuration only.

callbacks[6] = new

com.ibm.websphere.security.auth.callback.WSAppContextCallback("ApplicationContextCallback: ");

This callback handler collects the Web application context that is used during the login. This

callback handler consists of a HashMap object, which contains the application name and the redirect web address, if present. The callback handler is presented from the WEB_INBOUND login configuration only.

callbacks[7] = new WSRealmNameCallbackImpl("Realm Name: ", default_realm);

This callback handler collects the realm name for the login information. The realm information might not always be provided. If the realm information is not provided, assume that it is the current realm.

callbacks[8] = new WSX509CertificateChainCallback("X509Certificate[]: ");

This callback handler contains the certificate that was validated by Secure Sockets Layer (SSL) if the login source is an X509Certificate from SSL client authentication. The ItpaLoginModule calls the same mapping functions as WebSphere Application Server releases prior to version 6.1. However, having it passed into the login gives a custom login module the opportunity to map the certificate in a custom way. Then, it performs a Hashtable login. See "Configuring inbound identity mapping" on page 435 for more information on a Hashtable login.

- Use shared state variables to share information between login modules during the login phase.

If you want to access the objects that WebSphere Application Server creates during a login, refer to the following shared state variables. The variables are set in the following login modules: ItpaLoginModule, swamLoginModule, and wsMapDefaultInboundLoginModule.

Shared state variable

com.ibm.wsspi.security.auth.callback.Constants.WSPRINCIPAL_KEY

Purpose

Specifies the com.ibm.websphere.security.auth.WSPPrincipal object. See the WebSphere Application Server API documentation for application programming interface (API) usage. This shared state variable is for read-only purposes. Do not set this variable in the shared state for custom login modules.

The login module in which variables are set

ItpaLoginModule, swamLoginModule, and wsMapDefaultInboundLoginModule

Shared state variable

com.ibm.wsspi.security.auth.callback.Constants.WSCREDENTIAL_KEY

Purpose

Specifies the com.ibm.websphere.security.cred.WSCredential object. See the WebSphere Application Server API documentation for API usage. This shared state variable is for read-only purposes. Do not set this variable in the shared state for custom login modules.

Login module in which variables are set

wsMapDefaultInboundLoginModule

Shared state variable

com.ibm.wsspi.security.auth.callback.Constants.WSAUTHZTOKEN_KEY

Purpose

Specifies the default com.ibm.wsspi.security.token.AuthorizationToken object. Login modules can use this object to set custom attributes plugged in after the wsMapDefaultInboundLoginModule login module. The information set here is propagated downstream and is available to the application. See the WebSphere Application Server API documentation for API usage.

Initial versus propagation logins

As mentioned previously, some logins are considered initial logins because of the following reasons:

- It is the first time authentication information is presented to WebSphere Application Server.
- The login information is received from a server that does not propagate security attributes so this information must be gathered from a user registry.

Other logins are considered propagation logins when a `WSTokenHolderCallback` callback is present and contains sufficient information from a sending server to recreate all the required objects needed by WebSphere Application Server runtime. In cases where there is sufficient information for the WebSphere Application Server runtime, the information you might add to the Subject is likely to exist from the previous login. To verify if your object is present, you can get access to the `ArrayList` object that is present in the `WSTokenHolderCallback` callback, and search through this list looking at each `TokenHolder` `getName` method. This search is used to determine if WebSphere Application Server is deserializing your custom object during this login. Check the class name returned from the `getName` method using the `String` `startsWith` method because the runtime might add additional information at the end of the name to know which Subject is set to add the custom object after deserialization.

- Code your `login()` method to determine when sufficient information is present.

The following code snippet can be used in your `login()` method to determine when sufficient information is present. For another example, see “Configuring inbound identity mapping” on page 435.

```
// This is a hint provided by WebSphere Application Server that
// sufficient propagation information does not exist and, therefore,
// a login is required to provide the sufficient information. In this
// situation, a Hashtable login might be used.
boolean requiresLogin = ((com.ibm.wsspi.security.auth.callback.
WSTokenHolderCallback) callbacks[1]).requiresLogin();

if (requiresLogin)
{
// Check to see if your object exists in the TokenHolder list,
if not, add it.
java.util.ArrayList authzTokenList = ((WSTokenHolderCallback) callbacks[6]).
getTokenHolderList();boolean found = false;

if (authzTokenList != null)
{
Iterator tokenListIterator = authzTokenList.iterator();

while (tokenListIterator.hasNext())
{
com.ibm.wsspi.security.token.TokenHolder th = (com.ibm.wsspi.security.token.
TokenHolder) tokenListIterator.next();

if (th != null && th.getName().startsWith("com.acme.myCustomClass"))
{
found=true;
break;
}
}
if (!found)
{
// go ahead and add your custom object.
}
}
else
{
// This code indicates that sufficient propagation information is present.
// User registry calls are not needed by WebSphere Application Server to
// create a valid Subject. This code might be a no-op in your login module.
}
```

Sample custom login module

You can use the following sample to get ideas on how to use some of the callbacks and shared state variables.

```
{
// Defines your login module variables
com.ibm.wsspi.security.token.AuthenticationToken customAuthzToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthzToken = null;
com.ibm.websphere.security.cred.WSCredential credential = null;
```

```

com.ibm.websphere.security.auth.WSPincipal principal = null;
private javax.security.auth.Subject _subject;
private javax.security.auth.callback.CallbackHandler _callbackHandler;
private java.util.Map _sharedState;
private java.util.Map _options;

public void initialize(Subject subject, CallbackHandler callbackHandler,
    Map sharedState, Map options)
{
    _subject = subject;
    _callbackHandler = callbackHandler;
    _sharedState = sharedState;
    _options = options;
}

public boolean login() throws LoginException
{
    boolean succeeded = true;

    // Gets the CALLBACK information
    javax.security.auth.callback.Callback callbacks[] = new javax.security.
        auth.callback.Callback[7];
    callbacks[0] = new javax.security.auth.callback.NameCallback(
        "Username: ");
    callbacks[1] = new javax.security.auth.callback.PasswordCallback(
        "Password: ", false);
    callbacks[2] = new com.ibm.websphere.security.auth.callback.
        WSCredTokenCallbackImpl ("Credential Token: ");
    callbacks[3] = new com.ibm.wsspi.security.auth.callback.
        WSServletRequestCallback ("HttpServletRequest: ");
    callbacks[4] = new com.ibm.wsspi.security.auth.callback.
        WSServletResponseCallback ("HttpServletResponse: ");
    callbacks[5] = new com.ibm.wsspi.security.auth.callback.
        WSAppContextCallback ("ApplicationContextCallback: ");
    callbacks[6] = new com.ibm.wsspi.security.auth.callback.
        WSTokenHolderCallback ("Authz Token List: ");

    try
    {
        callbackHandler.handle(callbacks);
    }
    catch (Exception e)
    {
        // Handles exceptions
        throw new WSLoginFailedException (e.getMessage(), e);
    }

    // Sees which callbacks contain information
    uid = ((NameCallback) callbacks[0]).getName();
    char password[] = ((PasswordCallback) callbacks[1]).getPassword();
    byte[] credToken = ((WSCredTokenCallbackImpl) callbacks[2]).getCredToken();
    javax.servlet.http.HttpServletRequest request = ((WSServletRequestCallback)
        callbacks[3]).getHttpServletRequest();
    javax.servlet.http.HttpServletResponse response = ((WSServletResponseCallback)
        callbacks[4]).getHttpServletResponse();
    java.util.Map appContext = ((WSAppContextCallback)
        callbacks[5]).getContext();
    java.util.List authzTokenList = ((WSTokenHolderCallback)
        callbacks[6]).getTokenHolderList();

    // Gets the SHARED STATE information
    principal = (WSPincipal) _sharedState.get(com.ibm.wsspi.security.
        auth.callback.Constants.WSPRINCIPAL_KEY);
    credential = (WSCredential) _sharedState.get(com.ibm.wsspi.security.
        auth.callback.Constants.WSCREDENTIAL_KEY);
    defaultAuthzToken = (AuthorizationToken) _sharedState.get(com.ibm.
        wsspi.security.auth.callback.Constants.WSAUTHZTOKEN_KEY);
}

```



```

// What you tend to do with this information depends upon the scenario
// that you are trying to accomplish. This example demonstrates how to
// access various different information:
// - Determine if a login is initial versus propagation
// - Deserialize a custom authorization token (For more information, see
//   "Security attribute propagation" on page 448
// - Add a new custom authorization token (For more information, see
//   "Security attribute propagation" on page 448
// - Look for a WSCredential and read attributes, if found.
// - Look for a WSPincipal and read attributes, if found.
// - Look for a default AuthorizationToken and add attributes, if found.
// - Read the header attributes from the HttpServletRequest, if found.
// - Add an attribute to the HttpServletResponse, if found.
// - Get the Web application name from the appContext, if found.

// - Determines if a login is initial versus propagation. This is most
//   useful when login module is first.
boolean requiresLogin = ((WSTokenHolderCallback) callbacks[6]).requiresLogin();

// initial login - asserts privilege attributes based on user identity
if (requiresLogin)
{
    // If you are validating a token from another server, there is an
    // application programming interface (API) to get the uniqueID from it.
    if (credToken != null && uid == null)
    {
        try
        {
            String uniqueID = WSSecurityPropagationHelper.
                validateLTPAToken(credToken);
            String realm = WSSecurityPropagationHelper.getRealmFromUniqueID
                (uniqueID);
            // Now set it to the UID so you can use that to either map or
            // login with.
            uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueID);
        }
        catch (Exception e)
        {
            // handle exception
        }
    }

    // Adds a Hashtable to shared state.
    // Note: You can perform custom mapping on the NameCallback value returned
    // to change the identity based upon your own mapping rules.
    uid = mapUser (uid);

    // Gets the default InitialContext for this server.
    javax.naming.InitialContext ctx = new javax.naming.InitialContext();

    // Gets the local UserRegistry object.
    com.ibm.websphere.security.UserRegistry reg = (com.ibm.websphere.security.
        UserRegistry) ctx.lookup("UserRegistry");

    // Gets the user registry uniqueID based on the uid specified in the
    // NameCallback.
    String uniqueid = reg.getUniqueUserId(uid);
    uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueID);

    // Gets the display name from the user registry based on the uniqueID.
    String securityName = reg.getUserSecurityName(uid);

    // Gets the groups associated with this uniqueID.
    java.util.List groupList = reg.getUniqueGroupIds(uid);

    // Creates the java.util.Hashtable with the information you gathered from

```

```

    // the UserRegistry.
    java.util.Hashtable hashtable = new java.util.Hashtable();
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_UNIQUEID, uniqueid);
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_SECURITYNAME, securityName);
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_GROUPS, groupList);

    // Adds a cache key that is used as part of the lookup mechanism for
    // the created Subject. The cache key can be an Object, but should
    // implement the toString() method. Make sure the cacheKey contains
    // enough information to scope it to the user and any additional
    // attributes that you use. If you do not specify this property the
    // Subject is scoped to the WSCREDENTIAL_UNIQUEID returned, by default.
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_CACHE_KEY,
        "myCustomAttribute" + uniqueid);

    // Adds the hashtable to the sharedState of the Subject.
    _sharedState.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_PROPERTIES_KEY,hashtable);
}
// propagation login - process propagated tokens
else
{
    // - Deserializes a custom authorization token. For more information, see
    //   "Security attribute propagation" on page 448.
    //   This can be done at any login module plug in point (first,
    //   middle, or last).
    if (authzTokenList != null)
    {
        // Iterates through the list looking for your custom token
        for (int i=0; i<authzTokenList.size(); i++)
        {
            TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

            // Looks for the name and version of your custom AuthorizationToken
            // implementation
            if (tokenHolder.getName().equals("com.ibm.websphere.security.token.
                CustomAuthorizationTokenImpl") && tokenHolder.getVersion() == 1)
            {
                // Passes the bytes into your custom AuthorizationToken constructor
                // to deserialize
                customAuthzToken = new
                com.ibm.websphere.security.token.
                CustomAuthorizationTokenImpl(tokenHolder.getBytes());
            }
        }
    }
    // - Adds a new custom authorization token (For more information,
    //   see "Security attribute propagation" on page 448)
    //   This can be done at any login module plug in point (first, middle,
    //   or last).
}
else
{
    // Gets the PRINCIPAL from the default AuthenticationToken. This must
    // match all of the tokens.
    defaultAuthToken = (com.ibm.wsspi.security.token.AuthenticationToken)
        sharedState.get(com.ibm.wsspi.security.auth.callback.Constants.
            WSAUTHTOKEN_KEY);
    String principal = defaultAuthToken.getPrincipal();

    // Adds a new custom authorization token. This is an initial login.
    // Pass the principal into the constructor
    customAuthzToken = new com.ibm.websphere.security.token.

```

```

        CustomAuthorizationTokenImpl(principal);

// Adds any initial attributes
if (customAuthzToken != null)
{
    customAuthzToken.addAttribute("key1", "value1");
    customAuthzToken.addAttribute("key1", "value2");
    customAuthzToken.addAttribute("key2", "value1");
    customAuthzToken.addAttribute("key3", "something different");
}
}

// - Looks for a WSCredential and read attributes, if found.
// This is most useful when plugged in as the last login module.
if (credential != null)
{
    try
    {
        // Reads some data from the credential
        String securityName = credential.getSecurityName();
        java.util.ArrayList = credential.getGroupIds();
    }
    catch (Exception e)
    {
        // Handles exceptions
        throw new WSLoginFailedException (e.getMessage(), e);
    }
}

// - Looks for a WSPrincipal and read attributes, if found.
// This is most useful when plugged as the last login module.
if (principal != null)
{
    try
    {
        // Reads some data from the principal
        String principalName = principal.getName();
    }
    catch (Exception e)
    {
        // Handles exceptions
        throw new WSLoginFailedException (e.getMessage(), e);
    }
}

// - Looks for a default AuthorizationToken and add attributes, if found.
// This is most useful when plugged in as the last login module.
if (defaultAuthzToken != null)
{
    try
    {
        // Reads some data from the defaultAuthzToken
        String[] myCustomValue = defaultAuthzToken.getAttributes ("myKey");
        // Adds some data if not present in the defaultAuthzToken
        if (myCustomValue == null)
            defaultAuthzToken.addAttribute ("myKey", "myCustomData");
    }
    catch (Exception e)
    {
        // Handles exceptions
        throw new WSLoginFailedException (e.getMessage(), e);
    }
}

// - Reads the header attributes from the HttpServletRequest, if found.
// This can be done at any login module plug in point (first, middle,

```

```

    // or last).
    if (request != null)
    {
        java.util.Enumeration headerEnum = request.getHeaders();
        while (headerEnum.hasMoreElements())
        {
            System.out.println ("Header element: " + (String)headerEnum.nextElement());
        }
    }

    // - Adds an attribute to the HttpServletResponse, if found
    // This can be done at any login module plug in point (first, middle,
    // or last).
    if (response != null)
    {
        response.addHeader ("myKey", "myValue");
    }

    // - Gets the Web application name from the appContext, if found
    // This can be done at any login module plug in point (first, middle,
    // or last).
    if (appContext != null)
    {
        String appName = (String) appContext.get(com.ibm.wsspi.security.auth.
            callback.Constants.WEB_APP_NAME);
    }

    return succeeded;
}

public boolean commit() throws LoginException
{
    boolean succeeded = true;

    // Add any objects here that you have created and belong in the
    // Subject. Make sure the objects are not already added. If you added
    // any sharedState variables, remove them before you exit. If the abort()
    // method gets called, make sure you cleanup anything added to the
    // Subject here.

    if (customAuthzToken != null)
    {
        // Sets the customAuthzToken token into the Subject
        try
        {
            // Do this in a doPrivileged code block so that application code
            // does not need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Adds the custom authorization token if it is not
                        // null and not already in the Subject
                        if ((customAuthzTokenPriv != null) &&
                            (!_subject.getPrivateCredentials().contains(customAuthzTokenPriv)))
                        {
                            _subject.getPrivateCredentials().add(customAuthzTokenPriv);
                        }
                    }
                }
            });
        }
        catch (Exception e)
        {
            throw new WSLoginFailedException (e.getMessage(), e);
        }

        return null;
    }
}

```

```

    }
    });
}
catch (Exception e)
{
    throw new WSLoginFailedException (e.getMessage(), e);
}
}

return succeeded;
}

public boolean abort() throws LoginException
{
    boolean succeeded = true;

    // Makes sure to remove all objects that have already been added (both into the
    // Subject and shared state).

    if (customAuthzToken != null)
    {
        // remove the customAuthzToken token from the Subject
        try
        {
            final AuthorizationToken customAuthzTokenPriv = customAuthzToken;
            // Do this in a doPrivileged block so that application code does not need
            // to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Removes the custom authorization token if it is not
                        // null and not already in the Subject
                        if ((customAuthzTokenPriv != null) &&
                            (_subject.getPrivateCredentials().
                                contains(customAuthzTokenPriv)))
                        {
                            _subject.getPrivateCredentials().
                                remove(customAuthzTokenPriv);
                        }
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }
                }
            });
        }
        return null;
    }
    });
}
catch (Exception e)
{
    throw new WSLoginFailedException (e.getMessage(), e);
}
}

return succeeded;
}

public boolean logout() throws LoginException
{
    boolean succeeded = true;

    // Makes sure to remove all objects that have already been added

```

```

        // (both into the Subject and shared state).
    if (customAuthzToken != null)
    {
        // Removes the customAuthzToken token from the Subject
        try
        {
            final AuthorizationToken customAuthzTokenPriv = customAuthzToken;
            // Do this in a doPrivileged code block so that application code does
            // not need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.
                PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Removes the custom authorization token if it is not null and not
                        // already in the Subject
                        if ((customAuthzTokenPriv != null) && (_subject.
                            getPrivateCredentials().
                                contains(customAuthzTokenPriv)))
                        {
                            _subject.getPrivateCredentials().remove(customAuthzTokenPriv);
                        }
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }
                }
            });
        }
        catch (Exception e)
        {
            throw new WSLoginFailedException (e.getMessage(), e);
        }
    }

    return succeeded;
}
}

```

- Configure the system login for your custom login module.

Customizing application login with Java Authentication and Authorization Service

Using Java Authentication and Authorization Service (JAAS), you can customize your application login.

About this task

Java Authentication and Authorization Service (JAAS) is an API that enables applications to access authentication and access control services without being tied to those services. The following topics explaining customizing your application with JAAS are covered in this section:

- Developing programmatic logins with the Java Authentication and Authorization Service (JAAS)
- Configuring programmatic logins for JAAS
- Configuring a server-side Java Authentication and Authorization Service authentication and login configuration

Developing programmatic logins with the Java Authentication and Authorization Service

Use this topic to develop programmatic logins with the Java Authentication and Authorization Service.

Before you begin

Java Authentication and Authorization Service (JAAS) represents the strategic application programming interfaces (API) for authentication.

JAAS replaces the Common Object Request Broker Architecture (CORBA) programmatic login application programming interfaces (APIs).

WebSphere Application Server provides some extension to JAAS:

- Refer to the Developing applications that use CosNaming (CORBA Naming interface) article for details on how to set up the environment for thin client applications to access remote resources on a server.
- If the application uses a custom JAAS login configuration, verify that the JAAS login configuration is properly defined. See “Configuring programmatic logins for Java Authentication and Authorization Service” on page 403 for details.
- Some of the JAAS APIs are protected by Java 2 security permissions. If these APIs are used by application code, verify that these permissions are added to the application `was.policy` file.

For details, see the following articles:

- “Adding the `was.policy` file to applications” on page 756
- “Using PolicyTool to edit policy files” on page 743
- “Configuring the `was.policy` file” on page 750

For more details on which APIs are protected by Java 2 security permissions, check the IBM Developer Kit, Java Technology Edition; JAAS and WebSphere Application Server public APIs documentation in Security: Resources for learning.

Some of the APIs that are used in the sample code in this documentation and the Java 2 security permissions that are required by these APIs are in the following list:

- `javax.security.auth.login.LoginContext` constructors are protected by the `javax.security.auth.AuthPermission "createLoginContext"` object.
 - `javax.security.auth.Subject.doAs` and `com.ibm.websphere.security.auth.WSSubject.doAs` methods are protected by the `javax.security.auth.AuthPermission "doAs"` object.
 - `javax.security.auth.Subject.doAsPrivileged` and `com.ibm.websphere.security.auth.WSSubject.doAsPrivileged` methods are protected by the `javax.security.auth.AuthPermission "doAsPrivileged"` object.
- **Enhanced model to Java Platform, Enterprise Edition (Java EE) resources for authorization checks.**

Due to a design oversight in JAAS Version 1.0, the `javax.security.auth.Subject.getSubject` method does not return the Subject that is associated with the running thread inside a `java.security.AccessController.doPrivileged` code block. This oversight can present inconsistent behavior, which might have unwanted effects. The `com.ibm.websphere.security.auth.WSSubject` class provides a workaround to associate a Subject to a running thread. The `com.ibm.websphere.security.auth.WSSubject` class extends the JAAS model to Java Platform, Enterprise Edition (Java EE) resources for authorization checks. If the Subject associates with the running thread within the `com.ibm.websphere.security.auth.WSSubject.doAs` method or if the `com.ibm.websphere.security.auth.WSSubject.doAsPrivileged` code block contains product credentials, the Subject is used for Java EE resource authorization checks.

- **User interface support for defining new JAAS login configuration.**

You can configure a JAAS login configuration in the administrative console and store the JAAS login configuration in a configuration repository. Applications can define a new JAAS login configuration in the administrative console and the data is persisted in the configuration repository. However, WebSphere Application Server still supports the default JAAS login configuration format (plain text file) that is provided by the JAAS default implementation. If duplicate login configurations are defined in both the

configuration repository and the plain text file format, the one in the repository takes precedence.

Advantages to defining the login configuration in the configuration repository includes:

- Administrative console support in defining JAAS login configuration
- Central management of the JAAS login configuration
- Distribution of the JAAS login configuration

- **Application support for programmatic authentication.**

WebSphere Application Server provides JAAS login configurations for applications to perform programmatic authentication to the WebSphere security runtime. These configurations perform authentication to the WebSphere Application Server-configured authentication mechanism (Simple WebSphere Authentication Mechanism (SWAM) or Lightweight Third Party Authentication (LTPA)) and user registry (Local OS, Lightweight Directory Access Protocol (LDAP), custom registries, or federated repositories) based on the authentication data that is supplied. The authenticated Subject from these JAAS login configurations contains the required principal and credentials that the WebSphere security runtime can use to perform authorization checks on Java EE role-based protected resources.

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

Here are the JAAS login configurations that are provided by WebSphere Application Server:

- **WSLogin JAAS login configuration.** A generic JAAS login configuration can use Java clients, client container applications, servlets, JavaServer Pages (JSP) files, and Enterprise JavaBeans (EJB) components to perform authentication based on a user ID and password, or a token to the security runtime for WebSphere Application Server. However, this configuration does not honor the CallbackHandler handler that is specified in the client container deployment descriptor.
- **ClientContainer JAAS login configuration.** This JAAS login configuration honors the CallbackHandler handler that is specified in the client container deployment descriptor. The login module of this login configuration uses the CallbackHandler handler in the client container deployment descriptor if one is specified, even if the application code specified one callback handler in the login context. This is for a client container application.

A Subject authenticated with the previously mentioned JAAS login configurations contains a `com.ibm.websphere.security.auth.WSPPrincipal` principal and a `com.ibm.websphere.security.cred.WSCredential` credential. If the authenticated Subject is passed in the `com.ibm.websphere.security.auth.WSSubject.doAs` or the other `doAs` methods, the product security runtime can perform authorization checks on Java EE resources based on the `com.ibm.websphere.security.cred.WSCredential` Subject.

- **Customer-defined JAAS login configurations.**

You can define other JAAS login configurations to perform programmatic login which creates a custom Subject in either the client or server process. Certain credentials and principals are required in the Subject for the product security runtime to use it for sending authentication information from the client over a protocol or to use it for handling authorization on the server. The required credentials are generated from provided login modules.

The login module needed for a pure Java client login is as follows:

- `com.ibm.ws.security.common.auth.module.WSLoginModuleImpl` required;

In addition to using this login module, the callback handler used must be able to handle the following callback classes.

- `javax.security.auth.callback.NameCallback`
- `javax.security.auth.callback.PasswordCallback`

A username and password must be specified in the callback handler. Custom classes that are added to the Subject on the client side should get propagated to the server automatically whenever security attribute propagation is enabled. For information about enabling propagation for a pure Java client, see the corresponding step in “Propagating security attributes among application servers” on page 453.

Note: The classes added to the Subject must be Java serializable and de-serializable for this to occur properly.

The login modules needed for a server login are as follows:

- `com.ibm.ws.security.server.lm.ltpaLoginModule` required;
- `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` required;

For information about the callbacks used for a server-side login configuration, see “Customizing a server-side Java Authentication and Authorization Service authentication and login configuration” on page 425.

- **Naming requirements for programmatic login on a pure Java client.**

When programmatic login occurs on a pure Java client and the property `com.ibm.CORBA.validateBasicAuth` equals `true`, it is necessary for the security code to know where the `SecurityServer` resides. Typically, the default `InitialContext` is sufficient when a `java.naming.provider.url` property is set as a system property or when the property is set in the `jndi.properties` file. In other cases it is not desirable to have the same `java.naming.provider.url` properties set in a system-wide scope. In this case, there is a need to specify security specific bootstrap information in the `sas.client.props` file. The following steps present the order of precedence for determining how to find the `SecurityServer` in a pure Java client:

Note: Set `com.ibm.CORBA.validateBasicAuth=false` whenever connecting to a z/OS server. This function does not currently work from a distributed client to a z/OS server because the `SecurityServer` is located using the `UNAUTHENTICATED` principal, which is not accepted on a z/OS system.

1. Use the `sas.client.props` file and look for the following properties:

```
com.ibm.CORBA.securityServerHost=myhost.mydomain
com.ibm.CORBA.securityServerPort=mybootstrap port
```

If you specify these properties, you are guaranteed that security looks here for the `SecurityServer`. The host and port specified can represent any valid WebSphere host and bootstrap port. The `SecurityServer` resides on all server processes and therefore it is not important which host or port you choose. If specified, the security infrastructure within the client process look up the `SecurityServer` based on the information in the `sas.client.props` file.

2. Place the following code in your client application to get a new `InitialContext()`:

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...

// Perform an InitialContext and default lookup prior to logging
// in so that target realm and bootstrap host/port can be
// determined for SecurityServer lookup.

    Hashtable env = new Hashtable();
    env.put(Context.INITIAL_CONTEXT_FACTORY, "
        com.ibm.websphere.naming.WsnInitialContextFactory");
    env.put(Context.PROVIDER_URL,
        "corbaloc:iiop:myhost.mycompany.com:2809");
    Context initialContext = new InitialContext(env);
    Object obj = initialContext.lookup("");

// programmatic login code goes here.
```

Complete this step prior to running any programmatic login. It is in this code that you specify a URL provider for your naming context, but it must point to a valid WebSphere Application Server within the cell to which you are authenticating. Pointing to one cell allows thread specific programmatic logins going to different cells to have a single system-wide `SecurityServer` location.

3. Use the new default `InitialContext()` method relying on the naming precedence rules. These rules are defined in the article, Example: Getting the default initial context.

Example

See the “Example: Programmatic logins” article.

Example: Programmatic logins:

This example illustrates how application programs can perform a programmatic login using Java Authentication and Authorization Service (JAAS).

```
LoginContext lc = null;
```

```
try {
    lc = new LoginContext("WSLogin",
        new WSCallbackHandlerImpl("userName", "password"));
} catch (LoginException le) {
    System.out.println("Cannot create LoginContext. " + le.getMessage());
    // Insert the error processing code
} catch (SecurityException se) {
    System.out.println("Cannot create LoginContext." + se.getMessage());
    // Insert the error processing code
}

try {
    lc.login();
} catch (LoginException le) {
    System.out.println("Fails to create Subject. " + le.getMessage());
    // Insert the error processing code
}
```

As shown in the example, the new login context is initialized with the WSLogin login configuration and the WSCallbackHandlerImpl callback handler. Use the WSCallbackHandlerImpl instance on a server-side application where you do not want prompting. A WSCallbackHandlerImpl instance is initialized by the specified user ID, password, and realm information. The present WSLoginModuleImpl class implementation that is specified by the WSLogin login configuration can only retrieve authentication information from the specified callback handler. You can construct a login context with a Subject object, but the Subject is disregarded by the present WSLoginModuleImpl implementation. For product client-container applications, replace WSLogin login configuration by ClientContainer login configuration, which specifies the WSClientLoginModuleImpl implementation that is tailored for client container requirements.

For a pure Java application client, the product provides two other callback handler implementations: WSStdinCallbackHandlerImpl and WSGUICallbackHandlerImpl, which prompt for user ID, password, and realm information on the command line and pop-up panel, respectively. You can choose either of these product callback handler implementations, depending on the particular application environment. You can develop a new callback handler if neither of these implementations fit your particular application requirement.

You also can develop your own login module if the default WSLoginModuleImpl implementation fails to meet all your requirements. This product provides utility functions that the custom login module can use, which are described in the next section.

In cases where no java.naming.provider.url property is set as a system property or in the jndi.properties file, a default InitialContext context does not function if the product server is not at the localhost:2809 location. In this situation, construct a new InitialContext context programmatically ahead of the JAAS login. JAAS needs to know where the security server resides to verify that the entered user ID or password is correct, prior to performing a commit method. By constructing a new InitialContext context in the way specified below, the security code has the information that is needed to find the security server location and the target realm.

Note: The first line starting with `env.put` was split into two lines for illustration purposes only.

```
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...

// Perform an InitialContext and default lookup prior to logging in so that target realm
// and bootstrap host/port can be determined for SecurityServer lookup.

Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "corbaloc:iiop:myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
Object obj = initialContext.lookup("");

LoginContext lc = null;
try {
    lc = new LoginContext("WSLogin",
        new WSCallbackHandlerImpl("userName", "realm", "password"));
} catch (LoginException le) {
    System.out.println("Cannot create LoginContext. " + le.getMessage());
    // insert error processing code
} catch (SecurityException se) {
    System.out.println("Cannot create LoginContext." + se.getMessage());
    // Insert error processing
}

try {
    lc.login();
} catch (LoginException le) {
    System.out.println("Fails to create Subject. " + le.getMessage());
    // Insert error processing code
}
```

Programmatic login

Programmatic login is a type of form login that supports application presentation site-specific login forms for the purpose of authentication.

When enterprise bean client applications require the user to provide identifying information, the writer of the application must collect that information and authenticate the user. The work of the programmer can be broadly classified in terms of where the actual user authentication is performed:

- In a client program
- In a server program

Users of Web applications can receive prompts for authentication data in many ways. The `<login-config>` element in the Web application deployment descriptor file defines the mechanism that is used to collect this information. Programmers who want to customize login procedures, rather than relying on general purpose devices like a 401 dialog window in a browser, can use a form-based login to provide an application-specific HTML form for collecting login information.

No authentication occurs unless administrative security is enabled. If you want to use form-based login for Web applications, you must specify `FORM` in the `auth-method` tag of the `<login-config>` element in the deployment descriptor of each Web application.

Applications can present site-specific login forms by using the WebSphere Application Server form-login type. The Java Platform, Enterprise Edition (Java EE) specification defines form login as one of the authentication methods for Web applications. WebSphere Application Server provides a form-logout mechanism.

Java Authentication and Authorization Service programmatic login

Java Authentication and Authorization Service (JAAS) is a new feature in WebSphere Application Server. It is also mandated by the Java EE 1.4 Specification. JAAS is a collection of strategic authentication application programming interfaces (API) that replace the Common Object Request Broker Architecture (CORBA) programmatic login APIs. WebSphere Application Server provides some extensions to JAAS:

Before you begin developing with programmatic login APIs, consider the following points:

- For the pure Java client application or client container application, initialize the client Object Request Broker (ORB) security prior to performing a JAAS login. Do this by running the following code prior to the JAAS login:

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
// Perform an InitialContext and default lookup prior to logging
// in to initialize ORB security and for the bootstrap host/port
// to be determined for SecurityServer lookup. If you do not want
// to validate the userid/password during the JAAS login, disable
// the com.ibm.CORBA.validateBasicAuth property in the
// sas.client.props file.

Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
        "corbaloc:iiop:myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
Object obj = initialContext.lookup("");
```

For more information, see “Example: Programmatic logins” on page 392.

- For the pure Java client application or the client container application, make sure that the host name and the port number of the target Java Naming and Directory Interface (JNDI) bootstrap properties are specified properly. See the Developing applications that use CosNaming (CORBA Naming interface) section for details.
- If the application uses custom JAAS login configuration, make sure that the custom JAAS login configuration is properly defined. See the “Configuring programmatic logins for Java Authentication and Authorization Service” on page 403 section for details.
- Some of the JAAS APIs are protected by Java 2 security permissions. If these APIs are used by application code, make sure that these permissions are added to the application `was.policy` file. See “Adding the `was.policy` file to applications” on page 756, “Using PolicyTool to edit policy files” on page 743 and “Configuring the `was.policy` file” on page 750 sections for details. For more details of which APIs are protected by Java 2 Security permissions, check the IBM Developer Kit, Java Technology Edition; JAAS and the WebSphere Application Server public APIs documentation for more details. The following list contains the APIs that are used in the samples code provided in this documentation.
 - `javax.security.auth.login.LoginContext` constructors are protected by `javax.security.auth.AuthPermission "createLoginContext"`.
 - `javax.security.auth.Subject.doAs` and `com.ibm.websphere.security.auth.WSSubject.doAs` are protected by `javax.security.auth.AuthPermission "doAs"`.
 - `javax.security.auth.Subject.doAsPrivileged` and `com.ibm.websphere.security.auth.WSSubject.doAsPrivileged` are protected by `javax.security.auth.AuthPermission "doAsPrivileged"`.
- `com.ibm.websphere.security.auth.WSSubject`: Due to a design oversight in JAAS Version 1.0, `javax.security.auth.Subject.getSubject` does not return the Subject associated with the running thread inside a `java.security.AccessController.doPrivileged` code block. This can present an inconsistent behavior that is problematic and causes an undesirable effort to work around. The `com.ibm.websphere.security.auth.WSSubject` API provides a workaround to associate the Subject to the running thread. The `com.ibm.websphere.security.auth.WSSubject` API extends the JAAS model to Java

EE resources for authorization checks. The Subject that is associated with the running thread within `com.ibm.websphere.security.auth.WSSubject.doAs` or `com.ibm.websphere.security.auth.WSSubject.doAsPrivileged` code block is used for Java EE resources authorization checks.

- Administrative console support for defining new JAAS login configuration: You can configure JAAS login configuration in the administrative console and store it in the WebSphere Application Server configuration API. Applications can define new JAAS login configuration in the administrative console and the data is persisted in the configuration repository that is stored with the WebSphere Application Server configuration API. However, WebSphere Application Server still supports the default JAAS login configuration format that is provided by the JAAS default implementation. If duplication login configurations are defined in both the WebSphere Application Server configuration API and the plain text file format, the login configuration in the WebSphere Application Server configuration API takes precedence. Advantages to define the login configuration in the WebSphere Application Server configuration API include:
 - Defining the JAAS login configuration using the administrative console.
 - Managing the JAAS login configuration centrally.
 - Distributing the JAAS login configuration in a Network Deployment installation.
- JAAS login configurations for WebSphere Application Server: WebSphere Application Server provides JAAS login configurations for applications to perform programmatic authentication to the WebSphere Application Server security runtime. These JAAS login configurations for WebSphere Application Server perform authentication to the configured authentication mechanism, Simple WebSphere Authentication Mechanism (SWAM) or Lightweight Third-Party Authentication (LTPA), and user registry (Local OS, LDAP, or Custom) based on the authentication data supplied. The authenticated Subject from these JAAS login configurations contain the required principal and credentials that can be used by the WebSphere Application Server security runtime to perform authorization checks on Java EE role-based protected resources.

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

Here are the JAAS login configurations that are provided by WebSphere Application Server:

- *WSLogin JAAS login configuration:* A generic JAAS login configuration that a Java client, client container application, servlet, JSP file, enterprise bean, and so on, can use to perform authentication that is based on a user ID and password, or a token to the WebSphere Application Server security runtime. However, this configuration does not support the `CallbackHandler` handler that is specified in the client container deployment descriptor.
- *ClientContainer JAAS login configuration:* This JAAS login configuration recognizes the `CallbackHandler` handler that is specified in the client container deployment descriptor. The login module of this login configuration uses the `CallbackHandler` handler in the client container deployment descriptor if one is specified, even if the application code specified one `CallbackHandler` handler in the login context. This is for client container application.
- The Subjects that are authenticated with the previously mentioned JAAS login configurations contain a `com.ibm.websphere.security.auth.WSPrincipal` principal and a `com.ibm.websphere.security.auth.WSCredential` credential. If the authenticated Subject is passed to the `com.ibm.websphere.security.auth.WSSubject.doAs` method or the other `doAs` methods, the WebSphere Application Server security runtime can perform authorization checks on Java EE resources, based on the Subject `com.ibm.websphere.security.auth.WSCredential` credential.
- **Customer-defined JAAS login configurations:** You can define other JAAS login configurations. See “Configuring programmatic logins for Java Authentication and Authorization Service” on page 403 for details. Use these login configurations to perform programmatic authentication to the custom authentication mechanism. However, the subjects from these customer-defined JAAS login configurations might not be used by the WebSphere Application Server security runtime to perform authorization checks if the subject does not contain the required principal and credentials.

Finding the root cause login exception from a JAAS login

If you get a `LoginException` exception after issuing the `LoginContext.login` API, you can find the root cause exception from the configured user registry. In the login modules, the registry exceptions are wrapped by a `com.ibm.websphere.security.auth.WSLoginFailedException` class. This exception has a `getCause` method with which you can pull out the exception that was wrapped after issuing the previous command.

You are not always guaranteed to get a `WSLoginFailedException` exception, but most of the exceptions that are generated from the user registry display here. The following example illustrates a `LoginContext.login` API with the associated catch block. Cast the `WSLoginFailedException` exception to `com.ibm.websphere.security.auth.WSLoginFailedException` class if you want to issue the `getCause` API.

The following `determineCause` example can be used for processing `CustomUserRegistry` exception types.

```
try
{
    lc.login();
}
catch (LoginException le)
{
    // drill down through the exceptions as they might cascade through the runtime
    Throwable root_exception = determineCause(le);

    // now you can use "root_exception" to compare to a particular exception type
    // for example, if you have implemented a CustomUserRegistry type, you would
    // know what to look for here.
}

/* Method used to drill down into the WSLoginFailedException to find the
"root cause" exception */

public Throwable determineCause(Throwable e)
{
    Throwable root_exception = e, temp_exception = null;

    // keep looping until there are no more embedded WSLoginFailedException or
    // WSSecurityException exceptions
    while (true)
    {
        if (e instanceof com.ibm.websphere.security.auth.WSLoginFailedException)
        {
            temp_exception = ((com.ibm.websphere.security.auth.WSLoginFailedException)
                e).getCause();
        }
        else if (e instanceof com.ibm.websphere.security.WSSecurityException)
        {
            temp_exception = ((com.ibm.websphere.security.WSSecurityException)
                e).getCause();
        }
        else if (e instanceof javax.naming.NamingException)
            // check for Ldap embedded exception
            {
                temp_exception = ((javax.naming.NamingException)e).getRootCause();
            }
        else if (e instanceof your_custom_exception_here)
        {
            // your custom processing here, if necessary
        }
        else
        {
            // this exception is not one of the types we are looking for,
```



```

        // lets return now, this is the root from the WebSphere
        // Application Server perspective
        return root_exception;
    }
    if (temp_exception != null)
    {
        // we have an exception; go back and see if this has another
        // one embedded within it.
        root_exception = temp_exception;
        e = temp_exception;
        continue;
    }
    else
    {
        // we finally have the root exception from this call path, this
        // has to occur at some point
        return root_exception;
    }
}
}
}

```

Finding the root cause login exception from a Servlet filter

You can also receive the root cause exception from a servlet filter when addressing post-form login processing. This exception is useful because it shows the user what happened. You can issue the following API to obtain the root cause exception:

```

Throwable t = com.ibm.websphere.security.auth.WSSubject.getRootLoginException();
if (t != null)
    t = determineCause(t);

```

When you have the exception, you can run it through the previous determineCause example to get the native registry root cause.

Enabling root cause login exception propagation to pure Java clients

Currently, the root cause does not get propagated to a pure client for security reasons. However, you might want to propagate the root cause to a pure client in a trusted environment. If you want to enable root cause login exception propagation to a pure client, click **Security > Global security > Custom Properties** on the WebSphere Application Server Administrative Console and set the following property:

```
com.ibm.websphere.security.registry.propagateExceptionsToClient=true
```

Non-prompt programmatic login

WebSphere Application Server provides a non-prompt implementation of the `javax.security.auth.callback.CallbackHandler` interface, which is called `com.ibm.websphere.security.auth.callback.WSCallbackHandlerImpl`. Using this interface, an application can push authentication data to the WebSphere LoginModule instance to perform authentication. This capability is useful for server-side application code to authenticate an identity and to use that identity to invoke downstream Java EE resources.

```

javax.security.auth.login.LoginContext lc = null;

try {
    lc = new javax.security.auth.login.LoginContext("WSLogin",
        new com.ibm.websphere.security.auth.callback.WSCallbackHandlerImpl("user",
            "securityrealm", "securedpassword"));

    // create a LoginContext and specify a CallbackHandler implementation
    // CallbackHandler implementation determine how authentication data is collected
    // in this case, the authentication data is "push" to the authentication mechanism

```

```

// implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: failed to instantiate a LoginContext and the exception: "
+ e.getMessage());
e.printStackTrace();

// maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
// to the application, or the JAAS login configuration is not defined.
}

if (lc != null)
try {
lc.login(); // perform login
javax.security.auth.Subject s = lc.getSubject();
// get the authenticated subject

// Invoke a Java EE resource using the authenticated subject
com.ibm.websphere.security.auth.WSSubject.doAs(s,
new java.security.PrivilegedAction() {
public Object run() {
try {
bankAccount.deposit(100.00); // where bankAccount is a protected EJB
} catch (Exception e) {
System.out.println("ERROR: error while accessing EJB resource, exception: "
+ e.getMessage());
e.printStackTrace();
}
return null;
}
});
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: login failed with exception: " + e.getMessage());
e.printStackTrace();

// login failed, might want to provide relogin logic
}

```

You can use the `com.ibm.websphere.security.auth.callback.WSCallbackHandlerImpl` callback handler with a pure Java client, a client application container, enterprise bean, JavaServer Pages (JSP) files, servlet, or other Java 2 Platform, Enterprise Edition (Java EE) resources. See “Example: Programmatic logins” on page 392 for more information about Object Request Broker (ORB) security initialization requirements in a pure Java client.

Note: The `WSCallbackHandlerImpl` callback handler is different depending on whether you use WebSphere Application Server security or Web services security. It is located in the `sas.jar` file for security, and in the `was-wssecurity.jar` file for Web services security.

User interface prompt programmatic login

WebSphere Application Server also provides a user interface implementation of the `javax.security.auth.callback.CallbackHandler` implementation to collect authentication data from a user through user interface login prompts. The `com.ibm.websphere.security.auth.callback.WSGUICallbackHandlerImpl` callback handler presents a user interface login panel to prompt users for authentication data.

Note: This behavior requires an X11 server to be called out by the DISPLAY environment.

```

javax.security.auth.login.LoginContext lc = null;

try {

```

```

lc = new javax.security.auth.login.LoginContext("WSLogin",
new com.ibm.websphere.security.auth.callback.WSGUICallbackHandlerImpl());

// create a LoginContext and specify a CallbackHandler implementation
// CallbackHandler implementation determine how authentication data is collected
// in this case, the authentication data is collected by GUI login prompt
// and pass to the authentication mechanism implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: failed to instantiate a LoginContext and the exception: "
+ e.getMessage());
e.printStackTrace();

// maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
// to the application, or the JAAS login configuration is not defined.
}

if (lc != null)
try {
lc.login(); // perform login
javax.security.auth.Subject s = lc.getSubject();
// get the authenticated subject

// Invoke a Java EE resources using the authenticated subject
com.ibm.websphere.security.auth.WSSubject.doAs(s,
new java.security.PrivilegedAction() {
public Object run() {
try {
bankAccount.deposit(100.00); // where bankAccount is a protected enterprise bean
} catch (Exception e) {
System.out.println("ERROR: error while accessing EJB resource, exception: "
+ e.getMessage());
e.printStackTrace();
}
return null;
}
});
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: login failed with exception: " + e.getMessage());
e.printStackTrace();

// login failed, might want to provide relogin logic
}

```

Note: Do not use the `com.ibm.websphere.security.auth.callback.WSGUICallbackHandlerImpl` callback handler for server-side resources like enterprise bean, servlet, JSP files, and so on. The user interface login prompt blocks the server for user input. This behavior is not good for a server process.

WebSphere Application Server also provides a Kerberos credential cache implementation of the `javax.security.auth.callback.CallbackHandler` interface. The callback handler, `com.ibm.websphere.security.auth.callback.WSCcacheCallBackHandlerImpl`. Using this interface, an application can push authentication data to the WebSphere LoginModule instance to perform authentication.

This capability is only for the client side application code to authenticate to WebSphere Application Server with the Kerberos credential cache.

If the following options exist in the `wsjaas_client.conf` file, set them to false:

```

    useDefaultKeytab=false
    useDefaultCcache=false
    tryFirstPass=false
    useFirstPass=false
    forwardable=false
    renewable=false
    renewable=false
    noaddress=false
javax.security.auth.login.LoginContext lc = null;

String krb5Ccache = /etc/krb5/krb5cc_utt1e;

try {
    lc = new javax.security.auth.login.LoginContext("WSKRB5Login",
    new com.ibm.websphere.security.auth.callback.WSCcacheCallBackHandlerImpl(user, krb5Realm, krb5Ccache, false));
    // create a LoginContext and specify a CallbackHandler implementation
    // CallbackHandler implementation determines how authentication data is collected
    // in this case, the authentication data is collected by stdin prompt
    // and passed to the authentication mechanism implemented by the LoginModule.
    } catch (javax.security.auth.login.LoginException e) {
    System.err.println("ERROR: failed to instantiate a LoginContext and the exception:
        " + e.getMessage());
    e.printStackTrace();

    // maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
    // to the application, or the JAAS login configuration is not defined.
    }

    if (lc != null)
    try {
    lc.login(); // perform login
    javax.security.auth.Subject s = lc.getSubject();
    // get the authenticated subject

    // Invoke a Java EE resource using the authenticated subject
    com.ibm.websphere.security.auth.WSSubject.doAs(s,
    new java.security.PrivilegedAction() {
    public Object run() {
    try {
    bankAccount.deposit(100.00);
    // where bankAccount is a protected enterprise bean
    } catch (Exception e) {
    System.out.println("ERROR: error while accessing EJB resource, exception: "
        + e.getMessage());
    e.printStackTrace();
    }
    return null;
    }
    });
    } catch (javax.security.auth.login.LoginException e) {
    System.err.println("ERROR: login failed with exception: " + e.getMessage());
    e.printStackTrace();

    // login failed, might want to provide relogin logic
    }

```

Application Server with the default Kerberos credential cache.

```

javax.security.auth.login.LoginContext lc = null;

```

```

try {

```

```

lc = new javax.security.auth.login.LoginContext("WSKRB5Login",
new com.ibm.websphere.security.auth.callback.WSCacheCallbackHandlerImpl(user, krb5Realm, null, true));
// create a LoginContext and specify a CallbackHandler implementation
// CallbackHandler implementation determines how authentication data is collected
// in this case, the authentication data is collected by stdin prompt
// and passed to the authentication mechanism implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: failed to instantiate a LoginContext and the exception:
    " + e.getMessage());
e.printStackTrace();

// maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
// to the application, or the JAAS login configuration is not defined.
}

if (lc != null)
try {
lc.login(); // perform login
javax.security.auth.Subject s = lc.getSubject();
// get the authenticated subject

// Invoke a Java EE resource using the authenticated subject
com.ibm.websphere.security.auth.WSSubject.doAs(s,
new java.security.PrivilegedAction() {
public Object run() {
try {
bankAccount.deposit(100.00);
// where bankAccount is a protected enterprise bean
} catch (Exception e) {
System.out.println("ERROR: error while accessing EJB resource, exception: "
    + e.getMessage());
e.printStackTrace();
}
return null;
}
});
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: login failed with exception: " + e.getMessage());
e.printStackTrace();

// login failed, might want to provide relogin logic
}

```

Application Server with the Microsoft native Kerberos credential cache. The client must login to the Microsoft Domain Controller.

```

javax.security.auth.login.LoginContext lc = null;

try {
lc = new javax.security.auth.login.LoginContext("WSKRB5Login",
new com.ibm.websphere.security.auth.callback.WSCacheCallbackHandlerImpl(null, null, null, true));
// create a LoginContext and specify a CallbackHandler implementation
// CallbackHandler implementation determines how authentication data is collected
// in this case, the authentication data is collected by stdin prompt
// and passed to the authentication mechanism implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: failed to instantiate a LoginContext and the exception:
    " + e.getMessage());
e.printStackTrace();

// maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
// to the application, or the JAAS login configuration is not defined.

```

```

}

if (lc != null)
try {
lc.login(); // perform login
javax.security.auth.Subject s = lc.getSubject();
// get the authenticated subject

// Invoke a Java EE resource using the authenticated subject
com.ibm.websphere.security.auth.WSSubject.doAs(s,
new java.security.PrivilegedAction() {
public Object run() {
try {
bankAccount.deposit(100.00);
// where bankAccount is a protected enterprise bean
} catch (Exception e) {
System.out.println("ERROR: error while accessing EJB resource, exception: "
+ e.getMessage());
e.printStackTrace();
}
return null;
}
});
} catch (javax.security.auth.login.LoginException e) {
System.err.println("ERROR: login failed with exception: " + e.getMessage());
e.printStackTrace();

// login failed, might want to provide relogin logic
}

```

WSKRB5Login module

The WSKRB5Login JAAS login configuration: is a generic JAAS login configuration that a Java client, client container application, servlet, JSP file, or enterprise bean can use to perform authentication that is based on a Kerberos principal name password or a Kerberos credential cache to the WebSphere Application Server security runtime. However, this configuration does not support the CallbackHandler handler that is specified in the client container deployment descriptor.

Place either the **krb5.ini** or **krb5.conf** files you have created in a default location. If either file is not located in the default location you must set the `java.security.krb5.conf` JVM system property with the correct path and Kerberos configuration file name.

On a Windows[®] platform, the default location is `c:\winnt\krb5.ini`.

On a Linux[®] platform, the default location is `/etc/krb5.conf`.

On other Unix platforms, the default location is `/etc/krb5/krb5.conf`.

On a z/OS platform, the default location is `/etc/krb5/krb5.conf`.

Kerberos configuration settings, the Kerberos key distribution center (KDC) name, and realm settings are provided in the Kerberos configuration file or through `java.security.krb5.kdc` and `java.security.krb5.realm` system property files.

Stdin prompt programmatic login

WebSphere Application Server also provides a stdin implementation of the `javax.security.auth.callback.CallbackHandler` interface. The callback handler,

com.ibm.websphere.security.auth.callback.WSStdinCallbackHandlerImpl, prompts and collects authentication data from a user through the stdin prompt.

```
javax.security.auth.login.LoginContext lc = null;

try {
    lc = new javax.security.auth.login.LoginContext("WSLogin",
        new com.ibm.websphere.security.auth.callback.WSStdinCallbackHandlerImpl());

    // create a LoginContext and specify a CallbackHandler implementation
    // CallbackHandler implementation determines how authentication data is collected
    // in this case, the authentication data is collected by stdin prompt
    // and passed to the authentication mechanism implemented by the LoginModule.
} catch (javax.security.auth.login.LoginException e) {
    System.err.println("ERROR: failed to instantiate a LoginContext and the exception:
        " + e.getMessage());
    e.printStackTrace();

    // maybe javax.security.auth.AuthPermission "createLoginContext" is not granted
    // to the application, or the JAAS login configuration is not defined.
}

if (lc != null)
    try {
        lc.login(); // perform login
        javax.security.auth.Subject s = lc.getSubject();
        // get the authenticated subject

        // Invoke a Java EE resource using the authenticated subject
        com.ibm.websphere.security.auth.WSSubject.doAs(s,
            new java.security.PrivilegedAction() {
                public Object run() {
                    try {
                        bankAccount.deposit(100.00);
                        // where bankAccount is a protected enterprise bean
                    } catch (Exception e) {
                        System.out.println("ERROR: error while accessing EJB resource, exception: "
                            + e.getMessage());
                        e.printStackTrace();
                    }
                    return null;
                }
            });
    } catch (javax.security.auth.login.LoginException e) {
        System.err.println("ERROR: login failed with exception: " + e.getMessage());
        e.printStackTrace();

        // login failed, might want to provide relogin logic
    }
}
```

Note: Do not use the `com.ibm.websphere.security.auth.callback.WSStdinCallbackHandlerImpl` callback handler for server-side resources like enterprise beans, servlets, JSP files, and so on. The input from the stdin prompt is not sent to the server environment. Most servers run in the background and do not have a console. However, if the server does have a console, the stdin prompt blocks the server for user input. This behavior is not good for a server process.

Configuring programmatic logins for Java Authentication and Authorization Service

A new JAAS login configuration can be added and modified using the administrative console. The changes are saved in the cell-level security document and are available to all managed application servers.

Before you begin

Java Authentication and Authorization Service (JAAS) is a feature in WebSphere Application Server. JAAS is a collection of WebSphere Application Server strategic authentication APIs and replaces the Common Object Request Broker Architecture (CORBA) programmatic login APIs.

WebSphere Application Server provides some extensions to JAAS:

- **com.ibm.websphere.security.auth.WSSubject.** The `com.ibm.websphere.security.auth.WSSubject` API extends the JAAS authorization model to Java Platform, Enterprise Edition (Java EE) resources.
- You can configure the JAAS login in the administrative console and store this login configuration in the Application Server configuration. However, WebSphere Application Server still supports the default JAAS login configuration format (plain text file) that is provided by the JAAS default implementation. If duplicate login configurations are defined in both the WebSphere Application Server configuration API and the plain text file format, the one in the WebSphere Application Server configuration API takes precedence. Advantages to defining the login configuration in the WebSphere configuration API include:
 - User interface support in defining JAAS login configuration
 - Central management of the JAAS login configuration
 - Distribution of the JAAS login configuration during installation

Due to a design oversight in JAAS Version 1.0, the `javax.security.auth.Subject.getSubject` method does not return the subject that is associated with the running thread inside a `java.security.AccessController.doPrivileged` code block. This problem presents an inconsistent behavior that might cause unfavorable results. The `com.ibm.websphere.security.auth.WSSubject` API provides a workaround to associate the subject to a running thread.

- **Proxy LoginModule.** The Proxy LoginModule loads the actual LoginModule module. The default JAAS implementation does not use the thread context class loader to load classes. The LoginModule module cannot load if the LoginModule class file is not in the application class loader or the Java extension class loader class path. Due to this class loader visibility problem, WebSphere Application Server provides a proxy LoginModule module to load the JAAS LoginModule using the thread context class loader. You do not need to place the LoginModule implementation on the application class loader or the class path for the Java extension class loader with this proxy LoginModule module.

If you do not want to use the Proxy LoginModule module, you can place the LoginModule module in the `app_server_root/lib/ext/` directory. However, this action is not recommended due to the security risks.

JAAS login configurations are defined in the WebSphere Application Server configuration application programming interface (API) security document. Click **Security > Global security**. Under Java Authentication and Authorization Service, click **Application logins**. The following JAAS login configurations are available:

ClientContainer

Defines a login configuration and a LoginModule implementation that is similar to that of the WSLogin configuration, but enforces the requirements of the WebSphere Application Server client container. For more information, see “Configuration entry settings for Java Authentication and Authorization Service” on page 408.

DefaultPrincipalMapping,

Defines a special LoginModule module that is typically used by Java EE connectors to map an authenticated WebSphere Application Server user identity to a set of user authentication data (user ID and password) for the specified back-end enterprise information system (EIS). For more information about Java EE Connector and the DefaultMappingModule module, refer to the Java EE security section.

WSLogin

Defines a login configuration and a LoginModule implementation that applications can use in general.

A new JAAS login configuration can be added and modified using the administrative console. The changes are saved in the cell-level security document and are available to all managed application servers. An application server restart is required for the changes to take effect at run time.

Note: Do not remove or delete the predefined JAAS login configurations (such as, ClientContainer, WSLogin, and DefaultPrincipalMapping). Deleting or removing them can cause other enterprise applications to fail.

1. Delete a JAAS login configuration.
 - a. Click **Security > Global security**.
 - b. Under Java Authentication and Authorization Service, click **Application logins**. The Application Login Configuration panel is displayed.
 - c. Select the check box for the login configurations to delete and click **Delete**.
2. Create a new JAAS login configuration.
 - a. Click **Security > Global security**.
 - b. Under Java Authentication and Authorization Service, click **Application logins**.
 - c. Click **New**. The Application Login Configuration panel is displayed.
 - d. Specify the alias name of the new JAAS login configuration and click **Apply**. This value is the name of the login configuration that you pass in the `javax.security.auth.login.LoginContext` implementation for creating a new `LoginContext` context.
Click **Apply** to save changes and to add the extra node name that precedes the original alias name. Clicking **OK** does not save the new changes in the `security.xml` file.
 - e. Under Additional properties, click **JAAS Login Modules**.
 - f. Click **New**.
 - g. Specify the Module class name. Specify the WebSphere Application Server proxy `LoginModule` module because of the limitation of the class loader visibility.
 - h. Specify the `LoginModule` implementation as the `delegate` property of the Proxy `LoginModule` module. The WebSphere Application Server proxy `LoginModule` class name is `com.ibm.ws.security.common.auth.module.proxy.WSLoginModuleProxy`.
 - i. Select **Authentication strategy** from the list and click **Apply**.
 - j. Under Additional properties, click **Custom properties**. The Custom properties panel is displayed for the selected `LoginModule`.
 - k. Create a new property with the name `delegate` and the value of the real `LoginModule` implementation. You can specify other properties like `debug` with the `true` value. These properties are passed to the `LoginModule` class as options to the `initialize` method of the `LoginModule` instance.
 - l. Click **Save**.
For a Network Deployment installation, make sure that a file synchronization operation is performed to propagate the changes to other nodes.

Several locations are within the WebSphere Application Server directory structure where you can place a JAAS login module. The following list provides locations for the JAAS login module in order of recommendation:

- Within an enterprise archive (EAR) file for a specific Java Platform, Enterprise Edition (Java EE) application.

If you place the login module within the EAR file, the login module is accessible by the specific application only.

- In the WebSphere Application Server-shared library.

If you place the login module in the shared library, you must specify which applications can access the module. For more information on shared libraries, see *Managing shared libraries*.

- In the Java extensions directory.

If you place the JAAS login module in the Java extensions directory, the login module is available to all applications.

Although the Java extensions directory provides the greatest availability for the login module, place the login module in an application EAR file. If other applications need to access the same login module, consider using shared libraries.

3. Change the plain text file.

WebSphere Application Server supports the default JAAS login configuration format, which is a plain text file, that is provided by the JAAS default implementation. However, a tool is not provided that edits plain text files in this format. You can define the JAAS login configuration in the plain text file, which is located in the `app_server_root/properties/wsjaas.conf` file. Any syntax errors can cause the incorrect parsing of the plain JAAS login configuration text file. This problem can cause other applications to fail.

Java client programs that use the Java Authentication and Authorization Service (JAAS) for authentication must invoke with the JAAS configuration file specified. This configuration file is set in the `app_server_root/bin/launchClient.bat` file as:

```
set JAAS_LOGIN_CONFIG=-Djava.security.auth.login.config=%install_root%\properties\wsjaas_client.conf
```

If the **launchClient.bat** file is not used to invoke the Java client program, verify that the appropriate JAAS configuration file is passed to the Java virtual machine with the **-Djava.security.auth.login.config** flag.

Results

A new JAAS login configuration is created or an old JAAS login configuration is removed. An enterprise application can use a newly created JAAS login configuration without restarting the application server process.

However, new JAAS login configurations that are defined in the `app_server_root/properties/wsjaas.conf` file, do not refresh automatically. Restart the application servers to validate changes. These JAAS login configurations are specific to a particular node and are not available for other application servers running on other nodes.

What to do next

Create new JAAS login configurations that are used by enterprise applications to perform custom authentication. Use these newly defined JAAS login configurations to perform programmatic login.

Login configuration for Java Authentication and Authorization Service:

Java Authentication and Authorization Service (JAAS) is a new feature in WebSphere Application Server. JAAS is WebSphere Application Server strategic application programming interface (API) for authentication that replaces the Common Object Request Broker Architecture (CORBA) programmatic login API.

WebSphere Application Server provides some extensions to JAAS:

- **com.ibm.websphere.security.auth.WSSubject:** The `com.ibm.websphere.security.auth.WSSubject` API extends the JAAS authorization model to Java 2 Platform, Enterprise Edition (J2EE) resources. You can configure JAAS login in the administrative console or by using the scripting functions and store this configuration in the WebSphere Application Server configuration API. However, WebSphere Application Server still supports the default JAAS login configuration format, a plain text file, which is provided by the JAAS default implementation. If duplicate login configurations are defined in both the WebSphere Application Server configuration API and the plain text file format, the one in the WebSphere Application Server configuration API takes precedence. Advantages to defining the login configuration in the WebSphere configuration API include:
 - User interface support in defining JAAS login configuration
 - Central management of the JAAS login configuration
 - Distribution of the JAAS login configuration in a Network Deployment product installation

Due to a design oversight in JAAS 1.0, the `javax.security.auth.Subject.getSubject` method does not return the `Subject` that is associated with the running thread inside a `java.security.AccessController.doPrivileged` code block. This action can present an inconsistent behavior that is problematic. The `com.ibm.websphere.security.auth.WSSubject` extension provides a workaround to associate the `Subject` to the running thread. The `com.ibm.websphere.security.auth.WSSubject` extension expands the JAAS authorization model to J2EE resources.

Why WebSphere Application Server has its own subject class: You can retrieve the subjects in a `Subject.doAs` block with the `Subject.getSubject` call. However, this procedure does not work if an `AccessController.doPrivileged` call is contained within the `Subject.doAs` block. In the following example, `s1` is equal to `s`, but `s2` is null:

```
* AccessController.doPrivileged() not only truncates the Subject propagation,
* but also reduces the permissions. It does not include the JAAS security
* policy defined for the principals in the Subject.
Subject.doAs(s, new PrivilegedAction() {
    public Object run() {
        System.out.println("Within Subject.doAsPrivileged()");
        Subject s1 = Subject.getSubject(AccessController.getContext());
        AccessController.doPrivileged(new PrivilegedAction() {
            public Object run() {
                Subject s2 = Subject.getSubject(AccessController.getContext());
                return null;
            }
        });
        return null;
    }
});
```

- JAAS Login Configuration can be configured in either the administrative console or by using the scripting functions and stored in the WebSphere Application Server configuration repository. An application can define a new JAAS login configuration in the administrative console and persist the data in the configuration repository that is stored in the WebSphere Application Server configuration API. However, WebSphere Application Server still supports the default JAAS login configuration format that is provided by the JAAS default implementation. If duplicate login configurations are defined in both the WebSphere Application Server configuration API and the plan text file format, the one in the WebSphere Application Server configuration API takes precedence. The advantages to defining the login configuration in the WebSphere Application Server configuration API include:
 - UI support in defining JAAS login configuration.
 - The JAAS configuration login configuration can be managed centrally.
 - The JAAS configuration login configuration is distributed in a Network Deployment installation.
- **Proxy LoginModule:** The `Proxy.LoginModule` is a proxy to the configured user or the system-defined module that the context class loader uses to load the module instead of the system class loader. The default JAAS implementation does not use the thread context class loader to load classes. The `LoginModule` module cannot be loaded if the `LoginModule` class file is not in the application class loader or the class loader class path for the Java extension. WebSphere Application Server provides a proxy `LoginModule` module to load the JAAS `LoginModule` using the thread context class loader. You do not need to place the `LoginModule` implementation on the application class loader or the class loader class path for the Java extension with this proxy `LoginModule` module.

Note: Do not remove or delete the predefined JAAS login configurations (`ClientContainer`, `WSLogin` and `DefaultPrincipalMapping`). Deleting or removing them can cause other enterprise applications to fail.

A system administrator determines the authentication technologies, or login modules, to use for each application and configures them in a login configuration. The source of the configuration information, for example, a file or a database, is up to the current `javax.security.auth.login.Configuration` implementation. The WebSphere Application Server implementation permits the definition of the login configuration in both the WebSphere Application Server configuration API security document and in a JAAS configuration file, where the former takes precedence.

JAAS login configurations are defined in the API security document for WebSphere Application Server configuration for applications to use. To access the configurations, complete the following steps:

1. Click **Security > Global security**.
2. Under Java Authentication and Authorization Service, click **Application logins**.

The WSLgin module defines a login configuration and the LoginModule implementation that can be used by applications in general.

The ClientContainer module defines a login configuration and the LoginModule implementation that is similar to the WSLgin module, but enforces the requirements of the WebSphere Application Server client container.

The DefaultPrincipalMapping module defines a special LoginModule that is typically used by Java 2 Connector to map an authenticated WebSphere Application Server user identity to a set of user authentication data (user ID and password) for the specified back-end enterprise information system (EIS). For more information about Java 2 Connector and the DefaultMappingModule, see the Java 2 Security section.

A new JAAS login configuration can be added and modified using the administrative console. The changes are saved in the cell-level security document and are available to all managed application servers. An application server restart is required for the changes to take effect at runtime and for the client container login configuration to be made available.

WebSphere Application Server also reads JAAS configuration information from the `wsjaas.conf` file under the properties subdirectory of the root directory under which WebSphere Application Server is installed. Changes made to the `wsjaas.conf` file are used only by the local application server and take effect after the application server restarts. The JAAS configuration in the WebSphere Application Server configuration API security document takes precedence over that defined in the `wsjaas.conf` file. A configuration entry in the `wsjaas.conf` is overridden by an entry of the same alias name in the WebSphere Application Server configuration API security document.

The Java Authentication and Authorization Service (JAAS) login configuration entries in the administrative console are propagated to the server runtime when they are created, not when the configuration is saved. However, the deleted JAAS login configuration entries are not removed from the server runtime. To remove the entries, save the new configuration, then stop and restart the server.

The Samples Gallery provides a JAAS login sample that demonstrates how to use JAAS with WebSphere Application Server. The sample uses a server-side login with JAAS to authenticate a user with the security runtime for WebSphere Application Server. The sample demonstrates the following technology:

- Java 2 Platform, Enterprise Edition (J2EE) Java Authentication and Authorization Service (JAAS)
- JAAS for WebSphere Application Server
- WebSphere Application Server security

The form login sample is a component of the technology samples. For more information on how to access the form login sample, see [Accessing the Samples \(Samples Gallery\)](#).

Configuration entry settings for Java Authentication and Authorization Service:

Use this page to specify a list of Java Authentication and Authorization Service (JAAS) login configurations for the application code to use, including Java 2 Platform, Enterprise Edition (J2EE) components such as enterprise beans, JavaServer Pages (JSP) files, servlets, resource adapters, and message-driven beans (MDBs).

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > Application logins**.

Read the JAAS specifications before you begin defining additional login modules for authenticating to the application server security run time. You can define additional login configurations for your applications. However, if the application server LoginModule `com.ibm.ws.security.common.auth.module.WSLoginModuleImpl` module is not used or the LoginModule module does not produce a credential that is recognized by the application server. The application server security run time cannot use the authenticated subject from these login configurations for an authorization check for resource access.

You must invoke Java client programs that use Java Authentication and Authorization Service (JAAS) for authentication with a JAAS configuration file that is specified.

The application server supplies the `wsjaas_client.conf` default JAAS configuration file under the `app_server_root/properties` directory. This configuration file is set in the `app_server_root/bin/launchClient.bat` file as:

```
set JAAS_LOGIN_CONFIG=-Djava.security.auth.login.config=%app_server_root%\properties\wsjaas_client.conf
```

ClientContainer:

Specifies the login configuration used by the client container application, which uses the CallbackHandler API that is defined in the client container deployment descriptor.

The ClientContainer configuration is the default login configuration for the application server. Do not remove this default, as other applications that use it fail.

Default: ClientContainer

DefaultPrincipalMapping:

Specifies the login configuration that is used by Java 2 Connectors to map users to principals that are defined in the J2C authentication data entries.

The ClientContainer configuration is the default login configuration for the application server. Do not remove this default, as other applications that use it fail.

Default: ClientContainer

WSLogin:

Indicates whether all of the applications can use the WSLogin configuration to perform authentication for the application server security run time.

This login configuration does not honor the CallbackHandler handler that is defined in the client container deployment descriptor. To use this functionality, use the ClientContainer login configuration.

The WSLogin configuration is the default login configuration for the application server. Do not remove this default because other administrative applications that use it fail. This login configuration authenticates users for the application server security run time. Use the credentials from the authenticated subject that are returned from this login configuration as an authorization check for access to application server resources.

Default: ClientContainer

System login configuration entry settings for Java Authentication and Authorization Service:

Use this page to specify a list of Java Authentication and Authorization Service (JAAS) system login configurations.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > System logins**.

Read the Java Authentication and Authorization Service documentation before you begin defining additional login modules for authenticating to the application server security runtime. Do not remove the following system login modules:

- RMI_INBOUND
- WEB_INBOUND
- DEFAULT
- RMI_OUTBOUND
- SWAM
- wssecurity.IDAssertion
- wssecurity.signature
- wssecurity.PKCS7
- wssecurity.PkiPath
- wssecurity.UsernameToken
- wssecurity.X509BST
- LTPA
- LTPA_WEB

RMI_INBOUND, WEB_INBOUND, DEFAULT:

Processes inbound login requests for Remote Method Invocation (RMI), Web applications, and most of the other login protocols.

RMI_INBOUND

This login configuration handles logins for inbound RMI requests. Typically, these logins are requests for authenticated access to Enterprise JavaBeans (EJB) files. When using the RMI connector, these logins might be requests for Java Management Extensions (JMX).

WEB_INBOUND

This login configuration handles logins for Web application requests, which include servlets and JavaServer Pages (JSP) files. This login configuration can interact with the output that is generated from a trust association interceptor (TAI), if configured. The Subject that is passed into the WEB_INBOUND login configuration might contain objects that are generated by the TAI.

DEFAULT

This login configuration handles the logins for inbound requests that are made by most of the other protocols and internal authentications.

These three login configurations will pass in the following callback information, which is handled by the login modules within these configurations. These callbacks are not passed in at the same time. However, the combination of these callbacks determines how the application server authenticates the user.

Callback

```
callbacks[0] = new javax.security.auth.callback.  
NameCallback("Username: ");
```

Responsibility

Collects the user name that is provided during a login. This information can be the user name for the following types of logins:

- User name and password login, which is known as *basic authentication*.
- User name only for identity assertion.

Callback

```
callbacks[1] = new javax.security.auth.callback.  
PasswordCallback("Password: ", false);
```

Responsibility

Collects the password that is provided during a login.

Callback

```
callbacks[2] = new com.ibm.websphere.security.auth.callback.  
WSCredTokenCallbackImpl("Credential Token: ");
```

Responsibility

Collects the Lightweight Third Party Authentication (LTPA) token or other token type during a login. Typically, this information is present when a user name and a password are not present.

Callback

```
callbacks[3] = new com.ibm.wsspi.security.auth.callback.  
WSTokenHolderCallback("Authz Token List: ");
```

Responsibility

Collects the ArrayList list of the TokenHolder objects that are returned from the call to the WSOpaqueTokenHelper. The callback uses the createTokenHolderListFromOpaqueToken method with the Common Secure Interoperability version 2 (CSlv2) authorization token as input.

Note: This callback is present only when the **Security Attribute Propagation** option is enabled and this login is a propagation login. In a propagation login, sufficient security attributes are propagated with the request to prevent having to access the user registry for additional attributes. You must enable security attribute propagation for both the outbound and inbound authentication.

You can enable the **Security attribute propagation** option for CSlv2 outbound authentication by completing the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand RMI/IOP security and click **CSlv2 outbound authentication**.
3. Enable the **Security attribute propagation** option.

You can enable the **Security attribute propagation** option for CSlv2 inbound authentication by completing the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand RMI/IOP security and click **CSlv2 inbound authentication**.
3. Enable the **Security attribute propagation** option.

In system login configurations, the application server authenticates the user based upon the information that is collected by the callbacks. However, a custom login module does not need to act upon any of these callbacks. The following list explains the typical combinations of these callbacks:

- The `callbacks[0] = new javax.security.auth.callback.NameCallback("Username: ");` callback only
This callback occurs for CSlv2 identity assertion, Web and CSlv2 X509 certificate logins, old-style trust association interceptor logins, and so on. In Web and CSlv2 X509 certificate logins, the application server maps the certificate to a user name. This callback is used by any login type that establishes trust with the user name only.

- Both the `callbacks[0] = new javax.security.auth.callback.NameCallback("Username: ");` callback and the `callbacks[1] = new javax.security.auth.callback.PasswordCallback("Password: ", false);` callbacks.

This combination of callbacks is typical for basic authentication logins. Most user authentications occur using these two callbacks.

- The `callbacks[2] = new com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl("Credential Token: ");` only This callback is used to validate a Lightweight Third Party Authentication (LTPA) token. This validation typically occurs during a single sign-on (SSO) or downstream login. Any time a request originates from the application server, instead of a pure client, the LTPA token flows to the target server. For single sign-on (SSO), the LTPA token is received in the cookie and the token is used for login. If a custom login module needs the user name from an LTPA token, the module can use the following method to retrieve the unique ID from the token:

```
com.ibm.wsspi.security.token.WSSecurityPropagationHelper.  
validateLTPAToken(byte[])
```

After retrieving the unique ID, use the following method to get the user name:

```
com.ibm.wsspi.security.token.WSSecurityPropagationHelper.  
getUserFromUniqueID(uniqueID)
```

Note: Any time a custom login module is plugged in ahead of the application server login modules and it changes the identity using a credential mapping service, it is important that this login module validates the LTPA token, if present. Calling the following method is sufficient to validate the trust in the LTPA token:

```
com.ibm.wsspi.security.token.WSSecurityPropagationHelper.  
validateLTPAToken(byte[])
```

The receiving server must have the same LTPA keys as the sending server for this validation to be successful. A security exposure is possible if you do not validate this LTPA token, when present.

- A combination of any of the previously mentioned callbacks plus the `callbacks[3] = new com.ibm.wsspi.security.auth.callback.WSTokenHolderCallback("Authz Token List: ");` callback.

This callback indicates that some propagated attributes arrived at the server. The propagated attributes still require one of the following authentication methods:

- `callbacks[0] = new javax.security.auth.callback.NameCallback("Username: ");`
- `callbacks[1] = new javax.security.auth.callback.PasswordCallback("Password: ", false);`
- `callbacks[2] = new com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl("Credential Token: ");`

If the attributes are added to the Subject from a pure client, then the `NameCallback` and `PasswordCallback` callbacks authenticate the information and the objects that are serialized in the token holder are added to the authenticated Subject.

If both CSiv2 identity assertion and propagation are enabled, the application server uses the `NameCallback` callback and the token holder, which contains all of the propagated attributes, to deserialize most of the objects. The application server uses the `NameCallback` callback because trust is established with the servers that you indicate in the CSiv2 trusted server list. To specify trusted servers, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **CSiv2 inbound authentication**.

A custom login module needs to handle custom serialization. For more information, see "Security attribute propagation" in the information center.

In addition to the callbacks that are defined previously, the WEB_INBOUND login configuration can contain the following additional callbacks only:

Callback

```
callbacks[4] = new com.ibm.websphere.security.auth.callback.  
WSServletRequestCallback("HttpServletRequest: ");
```

Responsibility

Collects the HTTP servlet request object, if presented. This callback enables login modules to retrieve information from the HTTP request to use during a login.

Callback

```
callbacks[5] = new com.ibm.websphere.security.auth.callback.  
WSServletResponseCallback("HttpServletResponse: ");
```

Responsibility

Collects the HTTP servlet response object, if presented. This callback enables login modules to add information into the HTTP response as a result of the login. For example, login modules might add the SingleSignonCookie cookie to the response.

Callback

```
callbacks[6] = new com.ibm.websphere.security.auth.callback.  
WSApplicationContextCallback("ApplicationContextCallback: ");
```

Responsibility

Collects the Web application context used during the login. This callback consists of a hashtable, which if present contains the application name and the redirected Web address.

Callback

```
callbacks[7] = new WSRealmNameCallbackImpl("Realm Name: ", <default_realm>);
```

Responsibility

Collects the realm name for the login information. The realm information might not always be provided and should be assumed to be the current realm if it is not provided.

Callback

```
callbacks[8] = new WSX509CertificateChainCallback("X509Certificate[]: ");
```

Responsibility

If the login source is an X509Certificate from SSL client authentication, this callback contains the certificate that was validated by SSL. The ltpaLoginModule calls the same mapping functions as in previous releases. Once it is passed into the login, it provides a custom login module with the opportunity to map the certificate in a custom way. It then performs a hashtable login (see the related link, Custom login module for inbound mapping, for an example of a hashtable login.).

If you want to use security attribute propagation with the WEB_INBOUND login configuration, you can enable **Web inbound security attribute propagation** option on the Single sign-on panel.

1. Click **Security > Global security**.
2. Under Authentication, expand Web security and click **Single sign-on (SSO)**.
3. Select the **Web inbound security attribute propagation** option.

The following login modules are predefined for the RMI_INBOUND, WEB_INBOUND, and DEFAULT system login configurations. You can add custom login modules before, between, or after any of these login modules, but you cannot remove these predefined login modules:

- com.ibm.ws.security.server.lm.ltpaLoginModule

Performs the primary login when attribute propagation is either enabled or disabled. A primary login uses normal authentication information such as a user ID and password, an LTPA token, or a trust

association interceptor (TAI) and a certificate distinguished name (DN). If any of the following scenarios are true, this login module is not used and the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` module performs the primary login:

- The `java.util.Hashtable` object with the required user attributes is contained in the Subject.
 - The `java.util.Hashtable` object with the required user attributes is present in the `sharedState HashMap` of the `LoginContext`.
 - The `WSTokenHolderCallback` callback is present without a specified password. If a user name and a password are present with a `WSTokenHolderCallback` callback, which indicates propagated information, the request likely originates from either a pure client or a server from a different realm that mapped the existing identity to a user ID and password.
- `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule`

This login module performs the primary login using the normal authentication information if any of the following conditions are true:

- A `java.util.Hashtable` object with required user attributes is contained in the Subject.
- A `java.util.Hashtable` object with required user attributes is present in the `sharedState HashMap` of the `LoginContext` context.
- The `WSTokenHolderCallback` callback is present without a `PasswordCallback` callback.

When the `java.util.Hashtable` object is present, the login module maps the object attributes into a valid Subject. When the `WSTokenHolderCallback` callback is present, the login module deserializes the byte token objects and regenerates the serialized Subject contents. The `java.util.Hashtable` hashtable takes precedence over all of the other forms of login. Be careful to avoid duplicating or overriding what the application server might have propagated previously.

By specifying a `java.util.Hashtable` hashtable to take precedence over other authentication information, the custom login module must have already verified the LTPA token, if present, to establish sufficient trust. The custom login module can use the `com.ibm.wsspi.security.token.WSSecurityPropagationHelper.validationLTPAToken(byte[])` method to validate the LTPA token present in the `WSCredTokenCallback` callback. Failure to validate the LTPA token presents a security risk.

For more information on adding a hashtable containing well-known and well-formed attributes used by the application server as sufficient login information, see "Configuring inbound identity mapping" in the information center.

RMI_OUTBOUND:

Processes Remote Method Invocation (RMI) requests that are sent outbound to another server when either the `com.ibm.CSI.rmiOutboundLoginEnabled` or the `com.ibm.CSIOutboundPropagationEnabled` properties are true.

These properties are set in the CSiv2 authentication panel. To access the panel, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand RMI/IOP security and click **CSiv2 outbound authentication**.

To set the `com.ibm.CSI.rmiOutboundLoginEnabled` property, select **Custom outbound mapping**. To set the `com.ibm.CSIOutboundPropagationEnabled` property, select the **Security attribute propagation** option.

This login configuration determines the security capabilities of the target server and its security domain. For example, if the application server Version 5.1.1 or later (or 5.1.0.2 for z/OS) communicates with a Version 5.x Application Server, then the Version 5.1.1 Application Server sends the authentication information only, using an LTPA token, to the Version 5.x Application Server. However, if WebSphere Application Server Version 5.1.1 or later communicates with a Version 5.1.x Application Server, the authentication and authorization information is sent to the receiving application server if propagation is enabled at both the sending and receiving servers. When the application server sends both the

authentication and authorization information downstream, the application server removes the need to access the user registry again and look up the security attributes of the user for authorization purposes. Additionally, any custom objects that are added at the sending server are present in the Subject at the downstream server.

The following callback is available in the RMI_OUTBOUND login configuration. You can use the `com.ibm.wsspi.security.csiv2.CSiv2PerformPolicy` object that is returned by this callback to query the security policy for this particular outbound request. This query can help determine if the target realm is different than the current realm and if the application server must map the realm. For more information, see "Configuring outbound mapping to a different target realm" in the information center.

Callback

```
callbacks[0] = new WSProtocolPolicyCallback("Protocol Policy Callback: ");
```

Responsibility

Provides protocol-specific policy information for the login modules on this outbound invocation. This information is used to determine the level of security, including the target realm, target security requirements, and coalesced security requirements.

The following method obtains the `CSiv2PerformPolicy` policy from this specific login module:

```
csiv2PerformPolicy = (CSiv2PerformPolicy)
((WSProtocolPolicyCallback)callbacks[0]).getProtocolPolicy();
```

A different protocol other than RMI might have a different type of policy object.

The following login module is predefined in the RMI_OUTBOUND login configuration. You can add custom login modules before, between, or after any of these login modules, but you cannot remove these predefined login modules.

com.ibm.ws.security.Im.wsMapCSiv2OutboundLoginModule

Retrieves the following tokens and objects before creating an opaque byte that is sent to another server by using the Common Secure Interoperability Version 2 (CSiv2) authorization token layer:

- Forwardable `com.ibm.wsspi.security.token.Token` implementations from the Subject
- Serializable custom objects from the Subject
- Propagation tokens from the thread

You can use a custom login module prior to this login module to perform credential mapping. However, it is recommended that the login module change the contents of the Subject that is passed in during the login phase. If this recommendation is followed, the login modules are processed after this login module acts on the new Subject contents.

For more information, see "Configuring outbound mapping to a different target realm" in the information center.

SWAM:

Processes login requests in a single server environment when Simple WebSphere Authentication Mechanism (SWAM) is used as the authentication method.

SWAM does not support forwardable credentials. When SWAM is the authentication method, the application server cannot send requests from server to server. In this case, you must use LTPA.

Note: The SWAM login configuration is deprecated and will be removed in a future release.

wsssecurity.IDAssertion:

Version 5.x application

Processes login configuration requests for Web services security using identity assertion.

This login configuration is for Web Services Security Draft 13 JAX-RPC (Version 5.x) applications. For more information, see "Identity assertion authentication method" in the information center.

wssecurity.IDAssertionUsernameToken: **Version 6 and later applications**

Processes login configuration requests for Web services security using identity assertion.

This login configuration is for Web Services Security V1.0 JAX-RPC applications.

The custom property `com.ibm.wsspi.wssecurity.auth.module.IDAssertionLoginModule.disableUserRegistryCheck` can be configured for the JAAS `IDAssertionUsernameToken` login module. This property is an option for the Web services security identity assertion JAAS login module, `wssecurity.IDAssertionUsernameToken`. The property indicates that the login module should not perform a user registry check when processing an inbound identity token.

wssecurity.PKCS7:

Verifies an X.509 certificate with a certificate revocation list in a Public Key Cryptography Standards #7 (PKCS7) object.

This login configuration is for Version 6.0.x systems.

wssecurity.PkiPath:

Verifies an X.509 certificate with a public key infrastructure (PKI) path.

This login configuration is for Version 6.0.x systems.

wssecurity.signature:

Processes login configuration requests for Web services security using digital signature validation.

This login configuration is for Version 5.x systems.

wssecurity.UsernameToken:

Verifies basic authentication (user name and password).

When using the JAX-RPC runtime, the following custom properties can be configured for the JAAS `UsernameToken` login module:

The custom property `com.ibm.wsspi.wssecurity.auth.module.UsernameLoginModule.disableUserRegistryCheck` can be configured for the JAAS `UsernameToken` login module. This property is an option for the Web services security `UsernameToken` JAAS login module, `com.ibm.wsspi.wssecurity.auth.module.UsernameLoginModule`. The property indicates that the login module should not perform a user registry check when processing an inbound username token.

wssecurity.X509BST:

Verifies an X.509 binary security token (BST) by checking the validity of the certificate and the certificate path.

This login configuration is for Version 6.0.x systems.

LTPA_WEB:

Processes login requests to components in the Web container such as servlets and JavaServer pages (JSP) files.

The `com.ibm.ws.security.web.AuthenLoginModule` login module is predefined in the LTPA login configuration. You can add custom login modules before or after this module in the LTPA_WEB login configuration.

The LTPA_WEB login configuration can process the `HttpServletRequest` object, the `HttpServletResponse` object, and the Web application name that are passed in using a callback handler. For more information, see "Example: Customizing a server-side Java Authentication and Authorization Service authentication and logon configuration" in the information center.

LTPA:

Processes login requests that are not handled by the LTPA_WEB login configuration.

This login configuration is used by WebSphere Application Server Version 5.1 and previous versions.

The `com.ibm.ws.security.server.Im.ltpaLoginModule` login module is predefined in the LTPA login configuration. You can add custom login modules before or after this module in the LTPA login configuration. For more information, see "Example: Customizing a server-side Java Authentication and Authorization Service authentication and logon configuration" in the information center.

Login module settings for Java Authentication and Authorization Service:

Use this page to define the login module for a Java Authentication and Authorization Service (JAAS) login configuration.

You can define the JAAS login modules for application and system logins. To define these login modules in the administrative console, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > Application logins** or **System logins > alias_name**.
3. Under JAAS login modules, select a login module name to define.

Module class name:

Specifies the class name of the given login module.

Data type: String

Use login module proxy:

Specifies that the Java Authentication and Authorization Service (JAAS) loads the login module proxy class. JAAS then delegates calls to the login module classes that are defined in the Module class name field.

Use this option when you use both Version 5.x and Version 6 Application Servers in the same environment. If you migrate a Version 5.x Application Server to Version 6, WebSphere Application Server Version 6 automatically enables this option. If you have Version 6 only cells in your environment, you might choose to deselect this option.

Default: Enabled

Proxy class name:

Specifies the name of the proxy login module class.

The default login modules that are defined by the application server use the `com.ibm.ws.security.common.auth.module.WSLoginModuleProxy` proxy `LoginModule` class. This proxy class loads the application server login module with the thread context class loader and delegates all the operations to the real login module implementation. The real login module implementation is specified as the `delegate` option in the option configuration. The proxy class is needed because the Developer Kit application class loaders do not have visibility of the application server product class loaders.

Data type: String

Authentication strategy:

Specifies the authentication behavior as authentication proceeds down the list of login modules.

A Java Authentication and Authorization Service (JAAS) authentication provider supplies the authentication strategy. In JAAS, an authentication strategy is implemented through the `LoginModule` interface.

Data type: String
Default: Required
Range: Required, Requisite®, Sufficient and Optional

Required

The `LoginModule` module is required to succeed. Whether authentication succeeds or fails, the process still continues down the `LoginModule` list for each realm.

Requisite

The `LoginModule` module is required to succeed. If authentication is successful, the process continues down the `LoginModule` list in the realm entry. If authentication fails, control immediately returns to the application. Authentication does not proceed down the `LoginModule` list.

Sufficient

The `LoginModule` module is not required to succeed. If authentication succeeds, control immediately returns to the application. Authentication does not proceed down the `LoginModule` list. If authentication fails, the process continues down the list.

Optional

The `LoginModule` module is not required to succeed. Whether authentication succeeds or fails, the process still continues down the `LoginModule` list.

Specify additional options by clicking **Custom Properties** under Additional Properties. These name and value pairs are passed to the login modules during initialization. This process is one of the mechanisms that is used to pass information to login modules.

Module order:

Specifies the order in which the Java Authentication and Authorization Service (JAAS) login modules are processed.

Click **Set Order** to change the processing order of the login modules.

Login module order settings for Java Authentication and Authorization Service:

Use this page to specify the order in which the application server processes the login configuration modules.

You can specify the order of the login modules for application and system logins. To define these login modules in the administrative console, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > Application logins or System logins > alias**. You can create a new configuration by clicking **New**.
3. Under JAAS login modules, click **Set Order**.

When you select one of the JAAS login module class names, you can move that class name up and down the list. After you click **OK** and save the changes, the new order is reflected on either the Application login configuration or the System login configuration panel.

Login configuration settings for Java Authentication and Authorization Service:

Use this page to configure application login configurations.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **Java Authentication and Authorization Service > Application logins or System logins > alias_name**.

Click **Apply** to save changes and to add the extra node name that precedes the original alias name. Clicking **OK** does not save the new changes in the `security.xml` file.

Alias:

Specifies the alias name of the application login.

Do not use the forward slash character (/) in the alias name when defining JAAS login configuration entries. The JAAS login configuration parser cannot process the forward slash character.

Data type: String

Managing Java 2 Connector Architecture authentication data entries:

This task creates and deletes Java 2 Connector (J2C) authentication data entries.

Before you begin

Java Platform, Enterprise Edition (Java EE) Connector authentication data entries are used by resource adapters and Java DataBase Connectivity (JDBC) data sources. A Java EE Connector authentication data entry contains authentication data, which includes the following information:

Alias An identifier that identifies the authentication data entry. When configuring resource adapters or data sources, the administrator can specify which authentication data to choose using the corresponding alias.

User ID

A user identity of the intended security domain. For example, if a particular authentication data entry is used to open a new connection to DB2, this entry contains a DB2 user identity.

Password

The password of the user identity is encoded in the configuration repository.

Description

A short text description.

1. Delete a J2C authentication data entry.
 - a. Click **Security > Global security**.
 - b. Under Java Authentication and Authorization Service, click **J2C authentication data**. The **J2C Authentication Data Entries** panel is displayed.
 - c. Select the check boxes for the entries to delete and click **Delete**. Before deleting or removing an authentication data entry, make sure that it is not used or referenced by any resource adapter or data source. If the deleted authentication data entry is used or referenced by a resource, the application that uses the resource adapter or the data source fails to connect to the resources.
2. Create a new J2C authentication data entry.
 - a. Click **Security > Global security**.
 - b. Under Java Authentication and Authorization Service, click **J2C authentication data**. The **J2C Authentication Data Entries** panel is displayed.
 - c. Click **New**.
 - d. Enter a unique alias, a valid user ID, a valid password, and a short description (optional).

Note: When creating a new authentication data entry and specifying an alias, be aware that the node name is automatically appended as a prefix to the alias name you specify. For example, if you specify the alias, tek_test_Con3, and the node is MyNode, then the final alias name is: MyNode/tek_test_Con3.

Remember that this full alias name is the one used in your deployment descriptor.

- e. Click **OK** or **Apply**. No validation for the user ID and password is required.
- f. Click **Save**.

Make sure that a file synchronized operation is performed to propagate the changes to other nodes.

Results

A new J2C authentication data entry is created or an old entry is removed. The newly created entry is visible without restarting the application server process to use in the data source definition. But the entry is only in effect after the server is restarted. Specifically, the authentication data is loaded by an application server when starting an application and is shared among applications in the same application server.

If you create or update a data source that points to a newly created J2C authentication data alias, the test connection fails to connect until you restart the deployment manager. After you restart the deployment manager, the J2C authentication data is reflected in the run-time configuration. Any changes to the J2C authentication data fields require a deployment manager restart for the changes to take effect. The node agent must also be restarted.

What to do next

This step defines authentication data that you can share among resource adapters and data sources. Use the authentication data entry that is defined in the resource adapters or the data sources.

Java 2 Connector authentication data entry settings:

Use this page as a central place for administrators to define authentication data, which includes user identities and passwords. These values can reference authentication data entries by resource adapters, data sources, and other configurations that require authentication data using an alias.

You can display this page directly from the Java Authentication and Authorization Service (JAAS) configuration page or from other pages for resources that use J2EE Connector (J2C) authentication data entries. To view this administrative page, complete the following steps:

1. Click **Security > Global security**.

2. Under Authentication, click **Java Authentication and Authorization Service > J2C authentication data**.

Deleting authentication data entries: Be careful when deleting authentication data entries. If the deleted authentication data is used by other configurations, the initializing resources process fails.

Define a new authentication data entry by clicking **New**.

Alias:

Specifies the name of the authentication data entry.

Data type:	String
Units:	String
Default:	None

User ID:

Specifies the user identity.

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

Password:

Specifies the password that is associated with the user identity.

This field is not available on the collections panel. However, the panel is available when you create a new J2C authentication data entry.

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

Description:

Specifies an optional description of the authentication data entry. For example, this authentication data entry is used to connect to DB2.

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

J2C principal mapping modules:

You can develop your own J2EE Connector (J2C) mapping module if your application requires more sophisticated mapping functions. The mapping login module that you might have developed on WebSphere Application Server Version 5.x is still supported in WebSphere Application Server Version 6.0.x and later.

You can use the Version 5.x login modules in the connection factory mapping configuration. These login modules can also be used in the reference mapping configuration for the resource manager connection factory. A version 5.x mapping login module is not able to use the custom mapping properties.

If you want to develop a new mapping login module in Version 6.0.x and later, use the programming interface that is described in the following sections.

Note: Migrate your Version 5.x mapping login module to use the new programming model and the new custom properties as well as the mapping configuration isolation at application scope. Note that

mapping login modules that are developed using WebSphere Application Server Version 6.0.x cannot be used in the deprecated mapping configuration for the resource connection factory.

Invoking the login module for the resource reference mapping

A `com.ibm.wsspi.security.auth.callback.WSMappingCallbackHandler` class, which implements the `javax.security.auth.callback.CallbackHandler` interface, is a new WebSphere Application Service Provider Programming Interface (SPI) in WebSphere Application Server Version 6.0.x.

Application code uses the `com.ibm.wsspi.security.auth.callback.WSMappingCallbackHandlerFactory` helper class to retrieve a `CallbackHandler` object:

```
package com.ibm.wsspi.security.auth.callback;

public class WSMappingCallbackHandlerFactory {
    private WSMappingCallbackHandlerFactory;
    public static CallbackHandler getMappingCallbackHandler(
ManagedConnectionFactory mcf,
HashMap mappingProperties);
}
```

The `WSMappingCallbackHandler` class implements the `CallbackHandler` interface:

```
package com.ibm.wsspi.security.auth.callback;

public class WSMappingCallbackHandler implements CallbackHandler {
    public WSMappingCallbackHandler(ManagedConnectionFactory mcf,
HashMap mappingProperties);
    public void handle(Callback[] callbacks) throws IOException,
        UnsupportedCallbackException;
}
```

The `WSMappingCallbackHandler` handler can manage two new callback types that are defined in Version 6.0.x:

```
com.ibm.wsspi.security.auth.callback.WSManagedConnectionFactoryCallback
com.ibm.wsspi.security.auth.callback.WSMappingPropertiesCallback
```

The new login modules use the two callback types that are used at the reference mapping configuration for the resource manager connection factory. The `WSManagedConnectionFactoryCallback` callback provides a `ManagedConnectionFactory` instance that you set in the `PasswordCredential` credential. With this setting, the `ManagedConnectionFactory` instance can determine whether a `PasswordCredential` instance is used for signon to the target Enterprise Information Systems (EIS) instance. The `WSMappingPropertiesCallback` callback provides a hash map that contains custom mapping properties. The `com.ibm.mapping.authDataAlias` property name is reserved for setting the authentication data alias.

The WebSphere Application Server `WSMappingCallbackHandle` handle continues to support the two WebSphere Application Server Version 5.x callback types that older mapping login modules can use. The two callbacks defined can be used only by login modules that the login configuration uses at the connection factory. For backward compatibility, WebSphere Application Server Version 6.0.x and later passes the authentication data alias, if defined in the list of custom properties under the `com.ibm.mapping.authDataAlias` property name using the `WSAuthDataAliasCallback` callback to Version 5.x login modules:

```
com.ibm.ws.security.auth.j2c.WSManagedConnectionFactoryCallback
com.ibm.ws.security.auth.j2c.WSAuthDataAliasCallback
```

Invoking the login module for the connection factory mapping

The `WSPrincipalMappingCallbackHandler` class handles two callback types:

```
com.ibm.wsspi.security.auth.callback.WSManagedConnectionFactoryCallback
com.ibm.wsspi.security.auth.callback.WSMappingPropertiesCallback
```

The `WSPrincipalMappingCallbackHandler` handler and the two callbacks are deprecated in WebSphere Application Server Version 6.

Passing the mapping properties for the resource reference to the mapping login module

You can pass arbitrary custom properties to your mapping login module. The following example shows how the WebSphere Application Server default mapping login module looks for the authentication data alias property.

```
try {
    wspm_callbackHandler.handle(callbacks);
    String userID = null;
    String password = null;
    String alias = null;
    wspm_properties = ((WSMappingPropertiesCallback)callbacks[1]).getProperties();

    if (wspm_properties != null) {
        alias = (String) wspm_properties.get(com.ibm.wsspi.security.auth.callback.
            Constants.MAPPING_ALIAS);
        if (alias != null) {
            alias = alias.trim();
        }
    }
} catch (UnsupportedCallbackException unsupportedcallbackexception) {
    . . . // error handling
}
```

The default mapping login module for WebSphere Application Server Version 6.0.x requires one mapping property to define the authentication data alias. The mapping property, which is called `MAPPING_ALIAS`, is defined in the `Constants.class` file in the `com.ibm.wsspi.security.auth.callback` package.

```
MAPPING_ALIAS    =    "com.ibm.mapping.authDataAlias"
```

When you click **Use default method > Select authentication data entry authentication** on the Map resource references to resources panel, the administrative console automatically creates a `MAPPING_ALIAS` entry with the selected authentication data alias value in the mapping properties. If you create your own custom login configuration and then use the default mapping login module, you must set this property manually on the mapping properties for the resource factory reference.

In a custom login module, you can use the `WSSubject.getRunAsSubject` method to retrieve the subject that represents the identity of the current running thread. The identity of the current running thread is known as the *RunAs* identity. The *RunAs* subject typically contains a `WSPrincipal` principal in the principal set and a `WSCredential` credential in the public credential set. The subject instance that is created by your mapping module contains a `Principal` instance in the principals set and a `PasswordCredential` credential or an `org.ietf.jgss.GSSCredential` instance in the set of private credentials.

The `GenericCredential` interface that is defined in Java Cryptography Architecture (JCA) Specification Version 1.0 is removed in the JCA Version 1.5 specification. The `GenericCredential` interface is supported by WebSphere Application Server Version 6.0.x to support older resource adapters that might be programmed to the `GenericCredential` interface.

Customizing an application login to perform an identity assertion

Using the Java Authentication and Authorization Service (JAAS) login framework, you can create a JAAS login configuration that can be used to perform login to an identity assertion.

Before you begin

You can allow an application or system provider to perform an identity assertion with trust validation. To do this, you use the JAAS login framework, where trust validation is accomplished in one login module and

credential creation is accomplished in another module. The two custom login modules allow you to create a JAAS login configuration that can be used to perform a login to an identity assertion.

Two custom login modules are required:

User implemented trust association login module (trust validation)

The user implemented trust association login module performs whatever trust verification the user requires. When trust is verified, the trust verification status and the login identity should be put into a map in the share state of the login module so that the credential creation login module can use the information. This map should be stored in the property:

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.state`
(which consists of)

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.trusted`
(which is set to **true** if trusted and **false** if not trusted)

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.principal`
(which contains the principal of the identity)

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.certificates`
(which contains the certificate of the identity)

Identity assertion login module (credential creation)

The `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule` performs the credential creation. This module relies on the trust state information being in the login context's shared state. This login module is protected by the Java 2 security runtime permissions for:

- `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.initialize`
- `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.login`

The identity assertion login module looks for the trust information in the shared state property, `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.state`, which contains the trust status and the identity to login and should include:

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.trusted`
(which when **true** indicates trusted and **false** when not trusted)

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.principal`
(which contains the principal of the identity to login, if using a principal)

`com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.certificates`
(which contains a array of a certificate chain that contains the identity to login, if using a certificate)

A `WSLoginFailedException` is returned if the state, trust, or identity information is missing. The login module then performs a login of the identity, and the subject will contain the new identity

1. Delegate trust validation to a user implemented plug point. Trust validation must be accomplished in a custom login module. This custom login module should perform any trust validation required, then set the trust and identity information in the shared state to be passed on to the identity assertion login module. A map is required in the shared state key, `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.state`. If the state is missing then a `WSLoginFailedException` is thrown by the `IdentityAssertionLoginModule`. This map must include:
 - A trust key called `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.trust`. If the key is set to **true**, then trust is established. If the key is set to **false**, then no trust is established. If the trust key is not set to true, then the `IdentityAssertionLoginModule` will throw a `WSLoginFailedException`.

- An identity key is set: A java.security.Principal can be set in the com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.principal key.
- Or a java.security.cert.X509Certificate[] can be set in the com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.certificates key

If both a principal and certificate are supplied, then the principal is used and a warning is issued.

2. Create a new JAAS configuration for application logins The JAAS configuration will contain the user implemented trust validation custom login module and the IdentityAssertionLoginModule. Then to configure an application login configuration, perform the following on the administration console:
 - a. Expand **Security > Global security**.
 - b. Expand **Java authentication and authorization services > Application logins**
 - c. Select **New**.
 - d. Give the JAAS configuration an alias.
 - e. Click **Apply**.
 - f. Select **JAAS Login Modules**
 - g. Select **New**.
 - h. Enter the **Module class name** of the user implemented trust validation custom login module.
 - i. Click **Apply**.
 - j. Enter the **Module class name** of
com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule
 - k. Make sure the **Module class name** classes are in the correct order. The user implemented trust validation login module should be first and the IdentityAssertionLoginModule should be the second class in the list.
 - l. Click **Save**.

This JAAS configuration is then used by the application to perform an Identity Assertion.

3. Perform the programmable identity assertion. A program can now use the JAAS login configuration to perform a programmatic identity assertion. The application program can create a login context for the JAAS configuration created in step 2, then login to that login context with the identity they would assert to. If the login is successful then that identity can be set in the current running process. Here is an example of how such code would operate:

```
MyCallbackHandler handler = new MyCallbackHandler(new MyPrincipal("Joe"));
LoginContext lc = new LoginContext("MyAppLoginConfig", handler);
lc.login(); //assume successful
Subject s = lc.getSubject();
WSSubject.setRunAsSubject(s);
// From here on , the runas identity is "Joe"
```

Results

Using the JAAS login framework and two user implemented login modules, you can create a JAAS login configuration that can be used to perform login to an identity assertion.

Customizing a server-side Java Authentication and Authorization Service authentication and login configuration

WebSphere Application Server supports plugging in a custom Java Authentication and Authorization Service (JAAS) login module before or after the WebSphere Application Server system login module. However, WebSphere Application Server does not support the replacement of the WebSphere Application Server system login modules, which are used to create the WSCredential credential and WSPincipal principal in the Subject. By using a custom login module, you can either make additional authentication decisions or add information to the Subject to make additional, potentially finer-grained, authorization decisions inside a Java Platform, Enterprise Edition (Java EE) application.

About this task

WebSphere Application Server enables you to propagate information downstream that is added to the Subject by a custom login module. For more information, see “Security attribute propagation” on page 448. To determine which login configuration to use for plugging in your custom login modules, see the descriptions of the login configurations that are located in the “System login configuration entry settings for Java Authentication and Authorization Service” on page 409.

WebSphere Application Server supports the modification of the system login configuration through the administrative console and by using the wsadmin scripting utility. To configure the system login configuration using the administrative console, click **Security > Global security**. Under Java Authentication and Authorization Service, click **System logins**.

- Configure a system login configuration.

Refer to the following code sample to configure a system login configuration using the wsadmin tool. The following sample Jacl script adds a custom login module into the Lightweight Third-party Authentication (LTPA) Web system login configuration.

Note: Lines 32, 33, and 34 in the following code sample are split into two lines.

```
1. #####
2. #
3. # Open security.xml
4. #
5. #####
6.
7.
8. set sec [$AdminConfig getid /Cell:hillside/Security:/]
9.
10.
11. #####
12. #
13. # Locate systemLoginConfig
14. #
15. #####
16.
17.
18. set slc [$AdminConfig showAttribute $sec systemLoginConfig]
19.
20. set entries [lindex [$AdminConfig showAttribute $slc entries] 0]
21.
22.
23. #####
24. #
25. # Append a new LoginModule to LTPA_WEB
26. #
27. #####
28.
29. foreach entry $entries {
30. set alias [$AdminConfig showAttribute $entry alias]
31. if {$alias == "LTPA_WEB"} {
32. set newJAASLoginModuleId [$AdminConfig create JAASLoginModule
33. $entry {{moduleClassName
34. "com.ibm.ws.security.common.auth.module.proxy.WSLoginModuleProxy"}}]
35. set newPropertyId [$AdminConfig create Property $newJAASLoginModuleId
36. {{name delegate}{value "com.ABC.security.auth.CustomLoginModule"}}]
37. $AdminConfig modify $newJAASLoginModuleId
38. {{authenticationStrategy REQUIRED}}
39. break
40. }
41. }
```

```

40. #####
41. #
42. # save the change
43. #
44. #####
45.
46. $AdminConfig save 47.

```

Note: The wsadmin scripting utility inserts a new object to the end of the list. To insert the custom login module before the AuthenLoginModule login module, delete the AuthenLoginModule login module and recreate it after inserting the custom login module. Save the sample script into a sample.jacl file, and run the sample script using the following command:

```
wsadmin -f sample.jacl
```

- Remove the current LTPA_WEB login configuration and all of the login modules.

You can use the following sample Jacl script to remove the current LTPA_WEB login configuration and all the login modules:

```

48. #####
49. #
50. # Open security.xml
51. #
52. #####
53.
54.
55. set sec [$AdminConfig getid /Cell:hillside/Security:/]
56.
57.
58. #####
59. #
60. # Locate systemLoginConfig
61. #
62. #####
63.
64.
65. set slc [$AdminConfig showAttribute $sec systemLoginConfig]
66.
67. set entries [lindex [$AdminConfig showAttribute $slc entries] 0]
68.
69.
70. #####
71. #
72. # Remove the LTPA_WEB login configuration
73. #
74. #####
75.
76. foreach entry $entries {
77. set alias [$AdminConfig showAttribute $entry alias]
78. if {$alias == "LTPA_WEB"} {
79. $AdminConfig remove $entry
80. break 81. }
82. }
83.
84.
85. #####
86. #
87. # save the change
88. #
89. #####
90.
91. $AdminConfig save

```

- Recover the original LTPA_WEB configuration.

You can use the following sample Jacl script to recover the original LTPA_WEB configuration:

Note: Lines 122, 124, and 126 in the following code sample are split into two or more lines for illustrative purposes only.

```
92. #####
93. #
94. # Open security.xml
95. #
96. #####
97.
98.
99. set sec [$AdminConfig getid /Cell:hillside/Security:/]
100.
101.
102. #####
103. #
104. # Locate systemLoginConfig
105. #
106. #####
107.
108.
109. set slc [$AdminConfig showAttribute $sec systemLoginConfig]
110.
111. set entries [lindex [$AdminConfig showAttribute $slc entries] 0]
112.
113.
114.
115. #####
116. #
117. # Recreate the LTPA_WEB login configuration
118. #
119. #####
120.
121.
122. set newJAASConfigurationEntryId [$AdminConfig create
JAASConfigurationEntry $slc {{alias LTPA_WEB}}]
123.
124. set newJAASLoginModuleId [$AdminConfig create
JAASLoginModule $newJAASConfigurationEntryId
{moduleClassName
"com.ibm.ws.security.common.auth.module.proxy.WSLoginModuleProxy"}]
125.
126. set newPropertyId [$AdminConfig create Property
$newJAASLoginModuleId {{name delegate}{value "com.ibm.ws.security.web.AuthenLoginModule"}}]
127.
128. $AdminConfig modify $newJAASLoginModuleId
{{authenticationStrategy REQUIRED}}
129.
130.
131. #####
132. #
133. # save the change
134. #
135. #####
136.
137. $AdminConfig save
```

- Have the `ItpaLoginModule` initialize the callback array in the login method.

The WebSphere Application Server Version `ItpaLoginModule` and `AuthenLoginModule` login modules use the shared state to save state information so that custom login modules can modify the information. The `ItpaLoginModule` login module initializes the callback array in the login method using the following code. The callback array is created by the `ItpaLoginModule` login module only if an array is not defined in the shared state area. In the following code sample, the error handling code is removed to make the sample concise. If you insert a custom login module before the `ItpaLoginModule` login module, the custom login module might follow the same style to save the callback into the shared state.

Note: In the following code sample, several lines of code are split into two lines for illustrative purposes only.

```
138. Callback callbacks[] = null;
139. if (!sharedState.containsKey(
com.ibm.wsspi.security.auth.callback.Constants. CALLBACK_KEY)) {
140. callbacks = new Callback[3];
141. callbacks[0] = new NameCallback("Username: ");
142. callbacks[1] = new PasswordCallback("Password: ", false);
143. callbacks[2] = new com.ibm.websphere.security.auth.callback.
WSCredTokenCallbackImpl( "Credential Token: ");
144. try {
145. callbackHandler.handle(callbacks);
146. } catch (java.io.IOException e) {
147. . . .
148. } catch (UnsupportedCallbackException uce) {
149. . . .
150. }
151. sharedState.put( com.ibm.wsspi.security.auth.callback.
Constants.CALLBACK_KEY, callbacks);
152. } else {
153. callbacks = (Callback []) sharedState.get
( com.ibm.wsspi.security.auth.callback.Constants.CALLBACK_KEY);
154. }
```

- Have the AuthenLoginModule initialize the callback array.

The ItpaLoginModule and AuthenLoginModule login modules generate both a WSPPrincipal object and a WSCredential object to represent the authenticated user identity and security credentials. The WSPPrincipal and WSCredential objects also are saved in the shared state. A JAAS login uses a two-phase commit protocol.

First, the login methods in login modules, which are configured in the login configuration, are called. Then, their commit methods are called. A custom login module, which is inserted after the ItpaLoginModule and the AuthenLoginModule login modules, can modify the WSPPrincipal and WSCredential objects before these objects are committed. The WSCredential and WSPPrincipal objects must exist in the Subject after the login is completed. Without these objects in the Subject, WebSphere Application Server run-time code rejects the Subject to make security decisions.

AuthenLoginModule uses the following code to initialize the callback array:

Note: In the following code sample, several lines of code are split into two lines for illustrative purposes only.

```
155. Callback callbacks[] = null;
156. if (!sharedState.containsKey( com.ibm.wsspi.security.auth.
callback.Constants.CALLBACK_KEY)) {
157. callbacks = new Callback[6];
158. callbacks[0] = new NameCallback("Username: ");
159. callbacks[1] = new PasswordCallback("Password: ", false);
160. callbacks[2] = new com.ibm.websphere.security.auth.callback.
WSCredTokenCallbackImpl( "Credential Token: ");
161. callbacks[3] = new com.ibm.wsspi.security.auth.callback.
WSServletRequestCallback( "HttpServletRequest: ");
162. callbacks[4] = new com.ibm.wsspi.security.auth.callback.
WSServletResponseCallback( "HttpServletResponse: ");
163. callbacks[5] = new com.ibm.wsspi.security.auth.callback.
WSApplicationContextCallback( "ApplicationContextCallback: ");
164. try {
165. callbackHandler.handle(callbacks);
166. } catch (java.io.IOException e) {
167. . . .
168. } catch (UnsupportedCallbackException uce) {
169. . . .
170. }
171. sharedState.put( com.ibm.wsspi.security.auth.callback.
```

```

Constants.CALLBACK_KEY, callbacks);
172. } else {
173. callbacks = (Callback []) sharedState.get(com.ibm.wsspi.security.
auth.callback.Constants.CALLBACK_KEY);
174. }

```

- Obtain the application context. Three more objects, which contain callback information for the login, are passed from the Web container to the AuthenLoginModule login module: a `java.util.Map`, an `HttpServletRequest`, and an `HttpServletResponse` object. These objects represent the Web application context. The WebSphere Application Server Version 5.1 application context, `java.util.Map` object, contains the application name and the error page web address.

You can obtain the application context, `java.util.Map` object, by calling the `getContext` method on the `WSAppContextCallback` object. The `java.util.Map` object is created with the following deployment descriptor information.

Note: In the following code sample, several lines of code are split into two lines for illustrative purposes only.

```

175. HashMap appContext = new HashMap(2);
176. appContext.put( com.ibm.wsspi.security.auth.callback.
Constants.WEB_APP_NAME,web_application_name);
177. appContext.put( com.ibm.wsspi.security.auth.callback.Constants.
REDIRECT_URL,errorPage);

```

What to do next

The application name and the `HttpServletRequest` object might be read by the custom login module to perform mapping functions. The error page of the form-based login might be modified by a custom login module. In addition to the JAAS framework, WebSphere Application Server supports the trust association interface (TAI).

Other credential types and information can be added to the caller Subject during the authentication process using a custom login module. The third-party credentials in the caller Subject are managed by WebSphere Application Server as part of the security context. The caller Subject is bound to the running thread during the request processing. When a Web or an Enterprise JavaBeans (EJB) module is configured to use the caller identity, the user identity is propagated to the downstream service in an EJB request. The `WSCredential` credential and any third-party credentials in the caller Subject are not propagated downstream. Instead, some of the information can be regenerated at the target server based on the propagated identity. Add third-party credentials to the caller Subject at the authentication stage. The caller Subject, which is returned from the `WSSubject.getCallerSubject` method, is read-only and cannot be modified. For more information on the `WSSubject` subject, see "Getting the caller subject from the thread."

Getting the caller subject from the thread:

The Caller subject (or "received subject") contains the user authentication information that is used in the call for this request. This subject is returned after issuing the `WSSubject.getCallerSubject` application programming interface (API) to prevent replacing existing objects. The subject is marked read-only. This API can be used to get access to the `WSCredential` credential so that you can put or set data in the `hashmap` within the credential.

Before you begin

You need the following Java 2 security permissions to run this API: `permission javax.security.auth.AuthPermission "wssecurity.getCallerSubject;"`.

If you use the Kerberos authentication mechanism, the KDC policy enables Kerberos delegation and the client has a forwardable Kerberos ticket, the subject has the client Kerberos tickets and the GSS delegate credential. You can use APIs to access the Kerberos tickets and the GSS delegate credential.

About this task

Most data within the subject is not propagated downstream to another server. Only the credential token within the WSCredential credential is propagated downstream and a new caller subject is generated.

1. Get the caller subject.

```
caller_subject = com.ibm.websphere.security.auth.WSSubject.getCallerSubject();
```

2. Access the WSCredential credential.

```
caller_cred = caller_subject.getPublicCredentials(com.ibm.websphere
.security.cred.WSCredential.class).iterator().next();
```

3. Put or set data in the hashmap within the credential.

```
String CALLERDATA = (String) caller_cred.get ("MYKEY");
System.out.println("My data from the Caller credential is: " + CALLERDATA);
```

4. Access the Kerberos tickets. For example:

```
java.util.Set kerberosTickets = subject.getPrivateCredentials(KerberosTicket.class);
if ( kerberosTickets.size() > 1)
    System.out.println("Multiple Kerberos tickets found");
Iterator credIter = kerberosTickets.iterator();
```

5. Access the GSS credential. For example:

```
GSSCredential gssCred = subject.getPrivateCredentials(GSSCredential.class).iterator().next();
```

Example

```
try { javax.security.auth.Subject caller_subject; com.ibm.websphere.security.cred.WSCredential caller_cred
caller_subject = com.ibm.websphere.security.auth.WSSubject.getCallerSubject();
if (caller_subject != null) { caller_cred = caller_subject.getPublicCredentials
(com.ibm.websphere.security.cred.WSCredential.class).iterator().next();
String CALLERDATA = (String) caller_cred.get ("MYKEY");
System.out.println("My data from the Caller credential is: " + CALLERDATA); } }
catch (WSSecurityException e) { // log error } catch (Exception e) { // log error }
```

Related tasks

“Customizing application login with Java Authentication and Authorization Service” on page 388
Using Java Authentication and Authorization Service (JAAS), you can customize your application login.

Getting the RunAs subject from the thread:

The RunAs subject or invocation subject contains the user authentication information for the RunAs mode set in the application deployment descriptor for this method.

Before you begin

You need the Java 2 security permissions to run this API: permission javax.security.auth.AuthPermission "wssecurity.getRunAsSubject;".

About this task

The RunAs subject (or invocation subject) contains the user authentication information for the RunAs mode set in the application deployment descriptor for this method. This subject is marked read-only when returned from the WSSubject.getRunAsSubject application programming interface (API) to prevent replacing existing objects.

Most data within the Subject is not propagated downstream to another server. Only the credential token within the WSCredential credential is propagated downstream and a new Caller subject is generated.

1. Access the WSCredential credential. The WSCredential credential is documented in the API documentation.

2. Put or set data in the hashmap within the credential.

Example

```
try { javax.security.auth.Subject runas_subject; com.ibm.websphere.security.cred.WSCredential runas_cred;
runas_subject = com.ibm.websphere.security.auth.WSSubject.getRunAsSubject();
if (runas_subject != null) { runas_cred = runas_subject.getPublicCredentials
com.ibm.websphere.security.cred.WSCredential.class).iterator().next();
String RUNASDATA = (String) runas_cred.get ("MYKEY")
System.out.println("My data from the RunAs credential is: " + RUNASDATA ); } }
catch (WSSecurityException e) { // log error } catch (Exception e) { // log error }
```

Overriding the RunAs subject on the thread:

To extend the function that is provided by the Java Authentication and Authorization Service (JAAS) application programming interfaces (APIs), you can set the RunAs subject or invocation subject with a different valid entry that is used for outbound requests on this running thread.

Before you begin

You need the following Java 2 security permissions to run these APIs:

- permission javax.security.auth.AuthPermission "wssecurity.getRunAsSubject"
- permission javax.security.auth.AuthPermission "wssecurity.getCallerSubject"
- permission javax.security.auth.AuthPermission "wssecurity.setRunAsSubject"

About this task

This extension gives you the flexibility to associate the Subject with all the remote calls on this thread whether you use a WSSubject.doAs method to associate the subject with the remote action.

1. Set a new RunAs subject for the thread, overriding the one declaratively set.

```
com.ibm.websphere.security.auth.WSSubject.setRunAsSubject(caller_subject);
```

2. Perform some remote calls.

3. Restore the previous RunAs subject.

```
com.ibm.websphere.security.auth.WSSubject.setRunAsSubject(runas_subject);
```

Example

```
try { javax.security.auth.Subject runas_subject, caller_subject; runas_subject = com.ibm.websphere.security.
```

Revoking users from a cache:

In WebSphere Application Server, Version 5.0.2 and later, revocation of a user from the security cache using an MBean interface is supported.

About this task

When a user is removed from authentication cache, the user can still login to WebSphere Application Server at any time. Removing the cache only removes the user from the runtime cache. It does not remove the user from the registry, nor does it lock out the user.

Use the following procedure in a JAACL script.

The following Java Command Language (JAACL) revokes a user when given the realm and the user ID, and cycles through all the security administration MBean instances that are returned for the entire cell when run from the deployment manager wsadmin command. The command also purges the user from the cache during each process.

- e. Enter the module class name of the user-implemented trust validation custom login module, and then click **Apply**.
- f. Enter the `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule` module class name.
- g. Make sure that the module class name classes are in the correct order. The user-implemented trust validation login module must be the first class in the list, and the `IdentityAssertionLoginModule` module must be the second class.
- h. Click **Save**. The new JAAS configuration is used by the application to perform an identity assertion.

What to do next

An application can now use the JAAS login configuration to perform a programmatic identity assertion. The application can create a login context for the JAAS configuration created in step 2, then login to that login context with the identity it asserts to. If the login is successful, that identity can be set in the current running process, as in the following example:

```
MyCallbackHandler handler = new MyCallbackHandler(new MyPrincipal("Joe"));
LoginContext lc = new LoginContext("MyAppLoginConfig", handler);
lc.login(); //assume successful
Subject s = lc.getSubject();
WSSubject.setRunAsSubject(s);
// From here on, the runas identity is "Joe"
```

Performing identity mapping for authorization across servers in different realms

Identity mapping is a one-to-one mapping of a user identity between two servers so that the proper authorization decisions are made by downstream servers. Identity mapping is necessary when the integration of servers is needed, but the user registries are different and not shared between the systems.

About this task

In most cases, requests flow downstream between two servers that are part of the same security domain. In WebSphere Application Server, two servers that are members of the same cell are also members of the same security domain. In the same cell, the two servers have the same user registry and the same Lightweight Third Party Authentication (LTPA) keys for token encryption. These two commonalities ensure that the LTPA token, among other user attributes, which flows between the two servers, not only can be decrypted and validated, but also the user identity in the token can be mapped to attributes that are recognized by the authorization engine.

The most reliable and recommended configuration involves two servers within the same cell. However, sometimes you need to integrate multiple systems that cannot use the same user registry. When the user registries are different between two servers, the security domain or realm of the target server does not match the security domain of the sending server.

WebSphere Application Server enables mapping to occur either before sending the request outbound or before enabling the existing security credentials to flow to the target server. The credentials are mapped inbound with the specification that the target realm is trusted.

An alternative to mapping is to send the user identity without the token or the password to a target server without actually mapping the identity. The use of the user identity is based on trust between the two servers. Use Common Secure Interoperability Version 2 (CSIv2) identity assertion. When enabled, the server sends just the X.509 certificate, principal name, or distinguished name (DN) based upon what was used by the original client to perform the initial authentication. During CSIv2 identity assertion, trust is established between WebSphere Application Servers.

The user identity must exist in the target user registry for identity assertion to work. This process can also enable interoperability between other Java 2 Platform, Enterprise Edition (J2EE) Version 1.4 and higher compliant application servers. If both the sending server and target servers have identity assertion configured, WebSphere Application Server always uses this method of authentication, even when both servers are in the same security domain. For more information on CSiv2 identity assertion, see “Identity assertion to the downstream server” on page 503.

When the user identity is not present in the user registry of the target server, identity mapping must occur either before the request is sent outbound or when the request comes inbound. This decision depends upon your environment and requirements. However, it is typically easier to map the user identity before the request is sent outbound for the following reasons:

- You know the user identity of the existing credential as it comes from the user registry of the sending server.
- You do not have to worry about sharing Lightweight Third Party Authentication (LTPA) keys with the other target realm because you are not mapping the identity to LTPA credentials. Typically, you are mapping the identity to a user ID and password that are present in the user registry of the target realm.

When you do perform outbound mapping, in most cases, it is recommended that you use Secure Sockets Layer (SSL) to protect the integrity and confidentiality of the security information sent across the network. If LTPA keys are not shared between servers, an LTPA token cannot be validated at the inbound server. In this case, outbound mapping is necessary because the user identity cannot be determined at the inbound server to do inbound mapping. For more information, see “Configuring outbound mapping to a different target realm” on page 443.

When you need inbound mapping, potentially due to the mapping capabilities of the inbound server, you must ensure that both servers have the same LTPA keys so that you can get access to the user identity. Typically, in secure communications between servers, an LTPA token is passed into the WSCredTokenCallback callback of the inbound JAAS login configuration for the purposes of client authentication. A method is available that enables you to open the LTPA token, if valid, and get access to the user unique ID so that mapping can be performed. For more information, see “Configuring inbound identity mapping.” In other cases, such as identity assertion, you might receive a user name in the NameCallback callback of the inbound login configuration that enables you to map the identity.

The following topics are covered in this section:

- **Configuring inbound identity mapping** For inbound identity mapping, you can write a custom login module and configure WebSphere Application Server to run the login module first within the system login configurations. Consider the following steps when you write your custom login module: “Configuring inbound identity mapping.”
- **Configuring outbound identity mapping to a different target realm** By default, when WebSphere Application Server makes an outbound request from one server to another server in a different security realm, the request is rejected. This topic details alternatives for enabling one server to send outbound requests to a target server in a different realm. For more information, see “Configuring outbound mapping to a different target realm” on page 443

Configuring inbound identity mapping

For inbound identity mapping, write a custom login module and configure WebSphere Application Server to run the login module first within the system login configurations. Consider the following steps when you write your custom login module.

1. Get the inbound user identity from the callbacks and map the identity, if necessary This step occurs in the login method of the login module. A valid authentication has either or both NameCallback and the WSCredTokenCallback callbacks present. The following code sample shows you how to determine the user identity:

```
javax.security.auth.callback.Callback callbacks[] =
    new javax.security.auth.callback.Callback[3];
callbacks[0] = new javax.security.auth.callback.NameCallback("");
```

```

callbacks[1] = new javax.security.auth.callback.PasswordCallback
    ("Password: ", false);
callbacks[2] = new com.ibm.websphere.security.auth.callback.
    WSCredTokenCallbackImpl("");
callbacks[3] = new com.ibm.wsspi.security.auth.callback.
    WSTokenHolderCallback("");

try
{
    callbackHandler.handle(callbacks);
}
catch (Exception e)
{
    // Handles exceptions
    throw new WSLoginFailedException (e.getMessage(), e);
}

// Shows which callbacks contain information
boolean identitySwitched = false;
String uid = ((NameCallback) callbacks[0]).getName();
char password[] = ((PasswordCallback) callbacks[1]).getPassword();
byte[] credToken = ((WSCredTokenCallbackImpl) callbacks[2]).getCredToken();
java.util.List authzTokenList = ((WSTokenHolderCallback)
    callbacks[3]).getTokenHolderList();

if (credToken != null)
{
    try
    {
        String uniqueID = WSSecurityPropagationHelper.validateLTPAToken(credToken);
        String realm = WSSecurityPropagationHelper.getRealmFromUniqueID (uniqueID);
        // Now set the string to the UID so that you can use the result for either
        // mapping or logging in.
        uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueID);
    }
    catch (Exception e)
    {
        // Handles the exception
    }
}
else if (uid == null)
{
    // Throws an exception if authentication data is not valid.
    // You must have either UID or CredToken
    throw new WSLoginFailedException("invalid authentication data.");
}
else if (uid != null && password != null)
{
    // This is a typical authentication. You can choose to map this ID to
    // another ID or you can skip it and allow WebSphere Application Server
    // to log in for you. When passwords are presented, be very careful to not
    // validate the password because this is the initial authentication.

    return true;
}

// If desired, map this uid to something else and set the identitySwitched
// boolean. If the identity was changed, clear the propagated attributes
// below so they are not used incorrectly.
uid = myCustomMappingRoutine (uid);

// Clear the propagated attributes because they are no longer applicable
// to the new identity
if (identitySwitched)
{
    ((WSTokenHolderCallback) callbacks[3]).setTokenHolderList(null);
}

```

2. Check to see if attribute propagation occurred and if the attributes for the user are already present when the identity remains the same. Check to see if the user attributes are already present from the sending server to avoid duplicate calls to the user registry lookup. To check for the user attributes, use a method on the WSTokenHolderCallback callback that analyzes the information present in the callback to determine if the information is sufficient for WebSphere Application Server to create a Subject. The following code sample checks for the user attributes:

```
boolean requiresLogin =
((com.ibm.wsspi.security.auth.callback.WSTokenHolderCallback)
callbacks[2]).getrequiresLogin();
```

If sufficient attributes are not present to form the WSCredential and the WSPincipal objects that are needed to perform authorization, the previous code sample returns a true result. When the result is false, you can choose to discontinue processing as the necessary information exists to create the Subject without performing additional remote user registry calls.

3. Optional: Look up the required attributes from the user registry, put the attributes in a hashtable, and add the hashtable to the shared state. If the identity is switched in this login module, you must complete the following steps:
 - a. Create the hashtable of attributes, as shown in the following example.
 - b. Add the hashtable to the shared state.

If the identity is not switched, but the value of the requiresLogin code sample shown previously is true, you can create the hashtable of attributes. However, you are not required to create a hashtable in this situation as WebSphere Application Server handles the login for you. However, you might consider creating a hashtable to gather attributes in special cases where you are using your own special user registry. Creating a UserRegistry implementation, using a hashtable, and letting WebSphere Application Server gather the user attributes for you might be the easiest solution. The following table shows how to create a hashtable of user attributes:

```
if (requiresLogin || identitySwitched)
{
// Retrieves the default InitialContext for this server.
javax.naming.InitialContext ctx = new javax.naming.InitialContext();

// Retrieves the local UserRegistry implementation.
com.ibm.websphere.security.UserRegistry reg = (com.ibm.websphere.
security.UserRegistry)
ctx.lookup("UserRegistry");

// Retrieves the user registry uniqueID based on the uid specified
// in the NameCallback.
String uniqueid = reg.getUniqueUserId(uid);
uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueid);

// Retrieves the display name from the user registry based on the uniqueID.
String securityName = reg.getUserSecurityName(uid);

// Retrieves the groups associated with the uniqueID.
java.util.List groupList = reg.getUniqueGroupIds(uid);

// Creates the java.util.Hashtable with the information that you gathered
// from the UserRegistry implementation.
java.util.Hashtable hashtable = new java.util.Hashtable();
hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
WSCREDENTIAL_UNIQUEID, uniqueid);
hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
WSCREDENTIAL_SECURITYNAME, securityName);
hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
WSCREDENTIAL_GROUPS, groupList);

// Adds a cache key that is used as part of the lookup mechanism for
// the created Subject. The cache key can be an object, but should have
// an implemented toString method. Make sure that the cacheKey contains
```



```

    // enough information to scope it to the user and any additional attributes
    // that you are using. If you do not specify this property the Subject is
    // scoped to the returned WSCREDENTIAL_UNIQUEID, by default.
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_CACHE_KEY, "myCustomAttribute" + uniqueid);
    // Adds the hashtable to the sharedState of the Subject.
    _sharedState.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_PROPERTIES_KEY, hashtable);
}

```

The following rules define in more detail how a hashtable login is performed. You must use a `java.util.Hashtable` object in either the Subject (public or private credential set) or the shared-state `HashMap`. The `com.ibm.wsspi.security.token.AttributeNameConstants` class defines the keys that contain the user information. If the `Hashtable` object is put into the shared state of the login context using a custom login module that is listed prior to the Lightweight Third Party Authentication (LTPA) login module, the value of the `java.util.Hashtable` object is searched using the following key within the shared-state `HashMap`:

Property

`com.ibm.wsspi.security.cred.propertiesObject`

Reference to the property

`AttributeNameConstants.WSCREDENTIAL_PROPERTIES_KEY`

Explanation

This key searches for the `Hashtable` object that contains the required properties in the shared state of the login context.

Expected result

A `java.util.Hashtable` object.

If a `java.util.Hashtable` object is found either inside the Subject or within the shared state area, verify that the following properties are present in the hashtable:

Property

`com.ibm.wsspi.security.cred.uniqueId`

Reference to the property

`AttributeNameConstants.WSCREDENTIAL_UNIQUEID`

Returns

`java.util.String`

Explanation

The value of the property must be a unique representation of the user. For the WebSphere Application Server default implementation, this property represents the information that is stored in the application authorization table. The information is located in the application deployment descriptor after it is deployed and user-to-role mapping is performed. See the expected format examples if the user to role mapping is performed using a lookup to a WebSphere Application Server user registry implementation.

If a third-party authorization provider overrides the user-to-role mapping, then the third-party authorization provider defines the format. To ensure compatibility with the WebSphere Application Server default implementation for the unique ID value, call the WebSphere Application Server public `String` `getUniqueId(String userSecurityName)` `UserRegistry` method.

Expected format examples

Realm	Format (uniqueUserId)
Lightweight Directory Access Protocol (LDAP)	<code>ldaphost.austin.ibm.com:389/cn=user,o=ibm,c=us</code>
Windows	<code>MYWINHOST/S-1-5-21-963918322-163748893-4247568029-500</code>

Realm	Format (uniqueUserId)
UNIX	MYUNIXHOST/32

The com.ibm.wsspi.security.cred.uniqueId property is required.

Property

com.ibm.wsspi.security.cred.securityName

Reference to the property

AttributeNameConstants.WSCREDENTIAL_SECURITYNAME

Returns

java.util.String

Explanation

This property searches for the securityName of the authentication user. This name is commonly called the *display name* or *short name*. WebSphere Application Server uses the securityName attribute for the getRemoteUser, getUserPrincipal and getCallerPrincipal application programming interfaces (APIs). To ensure compatibility with the WebSphere Application Server default implementation for the securityName value, call the WebSphere Application Server public String getUserSecurityName(String uniqueUserId) UserRegistry method.

Expected format examples

Realm	Format (uniqueUserId)
LDAP	user (LDAP UID)
Windows	user (Windows username)
UNIX	user (UNIX username)

The com.ibm.wsspi.security.cred.securityName property is required.

Property

com.ibm.wsspi.security.cred.groups

Reference to the property

AttributeNameConstants.WSCREDENTIAL_GROUPS

Returns

java.util.ArrayList

Explanation

This key searches for the array list of groups to which the user belongs. The groups are specified in the *realm_name/user_name* format. The format of these groups is important as the groups are used by the WebSphere Application Server authorization engine for group-to-role mappings in the deployment descriptor. The format that is provided must match the format expected by the WebSphere Application Server default implementation. When you use a third-party authorization provider, you must use the format that is expected by the third-party provider. To ensure compatibility with the WebSphere Application Server default implementation for the unique group IDs value, call the WebSphere Application Server public List getUniqueGroupIds(String uniqueUserId) UserRegistry method.

Expected format examples for each group in the array list

Realm	Format
LDAP	ldap1.austin.ibm.com:389/cn=group1,o=ibm,c=us
Windows	MYWINREALM/S-1-5-32-544
UNIX	MY/S-1-5-32-544

The `com.ibm.wsspi.security.cred.groups` property is not required. A user is not required to have associated groups.

Property

`com.ibm.wsspi.security.cred.cacheKey`

Reference to the property

`AttributeNameConstants.WSCREDENTIAL_CACHE_KEY`

Returns

`java.lang.Object`

Explanation

This key property can specify an object that represents the unique properties of the login, including the user-specific information and the user dynamic attributes that might affect uniqueness. For example, when the user logs in from location A, which might affect their access control, the cache key needs to include location A so that the Subject that is received is the correct Subject for the current location.

This `com.ibm.wsspi.security.cred.cacheKey` property is not required. When this property is not specified, the cache lookup is the value that is specified for `WSCREDENTIAL_UNIQUEID`. When this information is found in the `java.util.Hashtable` object, WebSphere Application Server creates a Subject similar to the Subject that goes through the normal login process at least for LTPA. The new Subject contains a `WSCredential` object and a `WSPrincipal` object that is fully populated with the information found in the `Hashtable` object.

4. Add your custom login module into the `RMI_INBOUND`, `WEB_INBOUND`, and `DEFAULT` Java Authentication and Authorization Service (JAAS) system login configurations. Configure the `RMI_INBOUND` login configuration so that WebSphere Application Server loads your new custom login module first.
 - a. Click **Security > Global security > Java Authentication and Authorization Service > System logins > RMI_INBOUND**
 - b. Under Additional Properties, click **JAAS login modules > New** to add your login module to the `RMI_INBOUND` configuration.
 - c. Return to the JAAS login modules panel for `RMI_INBOUND`.
 - d. Click **Set order** to change the order that the login modules are loaded so that WebSphere Application Server loads your custom login module first. Use the **Move Up** or **Move Down** buttons to arrange the order of the login modules.
 - e. Repeat the previous three steps for the `WEB_INBOUND` and `DEFAULT` login configurations.

Results

This process configures identity mapping for an inbound request.

Example

The “Example: Custom login module for inbound mapping” topic shows a custom login module that creates a `java.util.Hashtable` hashtable that is based on the specified `NameCallback` callback. The `java.util.Hashtable` hashtable is added to the `sharedState` `java.util.Map` map so that the WebSphere Application Server login modules can locate the information in the hashtable.

Example: Custom login module for inbound mapping

This sample shows a custom login module that creates a `java.util.Hashtable` hashtable that is based on the specified `NameCallback` callback. The `java.util.Hashtable` hashtable is added to the `sharedState` `java.util.Map` map so that the WebSphere Application Server login modules can locate the information in the `Hashtable`.

```

public customLoginModule()
{

public void initialize(Subject subject, CallbackHandler callbackHandler,
    Map sharedState, Map options)
{
    // (For more information on initialization, see
    // "Developing custom login modules for a system login configuration" on page 375.)
    _sharedState = sharedState;
}

public boolean login() throws LoginException
{
    // (For more information on what to do during login, see
    // "Developing custom login modules for a system login configuration" on page 375.)

    // Handles the WSTokenHolderCallback to see if this is an initial or
    // propagation login.
    javax.security.auth.callback.Callback callbacks[] =
        new javax.security.auth.callback.Callback[3];
    callbacks[0] = new javax.security.auth.callback.NameCallback("");
    callbacks[1] = new javax.security.auth.callback.PasswordCallback(
        "Password: ", false);
    callbacks[2] = new com.ibm.websphere.security.auth.callback.
        WSCredTokenCallbackImpl("");
    callbacks[3] = new com.ibm.wsspi.security.auth.callback.
        WSTokenHolderCallback("");

    try
    {
        callbackHandler.handle(callbacks);
    }
    catch (Exception e)
    {
        // Handles the exception
    }

    // Determines which callbacks contain information
    boolean identitySwitched = false;
    String uid = ((NameCallback) callbacks[0]).getName();
    char password[] = ((PasswordCallback) callbacks[1]).getPassword();
    byte[] credToken = ((WSCredTokenCallbackImpl) callbacks[2]).getCredToken();
    java.util.List authzTokenList = ((WSTokenHolderCallback) callbacks[3]).
        getTokenHolderList();

    if (credToken != null)
    {
        try
        {
            String uniqueID = WSSecurityPropagationHelper.validateLTPAToken(credToken);
            String realm = WSSecurityPropagationHelper.getRealmFromUniqueID (uniqueID);
            // Set the string to the UID so you can use the information to either
            // map or login.
            uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueid);
        }
        catch (Exception e)
        {
            // handle exception
        }
    }
    else if (uid == null)
    {
        // The authentication data is not valid. You must have either UID
        // or CredToken
        throw new WSLoginFailedException("invalid authentication data.");
    }
    else if (uid != null && password != null)

```

```

{
    // This is a typical authentication. You can choose to map this ID to
    // another ID or you can skip it and allow WebSphere Application Server
    // to log in for you. When passwords are presented, be very careful not
    // to validate the password because this is the initial authentication.

    return true;
}

// You can map this uid to something else and set the identitySwitched
// boolean. If the identity is changed, clear the following propagated
// attributes so they are not used incorrectly.
uid = myCustomMappingRoutine (uid);

// Clear the propagated attributes because they no longer apply to the new identity
if (identitySwitched)
{
    ((WSTokenHolderCallback) callbacks[3]).setTokenHolderList(null);
}
boolean requiresLogin = ((com.ibm.wsspi.security.auth.callback.
    WSTokenHolderCallback) callbacks[2]).getRequiresLogin();

if (requiresLogin || identitySwitched)
{
    // Retrieves the default InitialContext for this server.
    javax.naming.InitialContext ctx = new javax.naming.InitialContext();

    // Retrieves the local UserRegistry object.
    com.ibm.websphere.security.UserRegistry reg =
        (com.ibm.websphere.security.UserRegistry) ctx.lookup("UserRegistry");

    // Retrieves the registry uniqueID based on the uid that is specified
    // in the NameCallback.
    String uniqueid = reg.getUniqueUserId(uid);
    uid = WSSecurityPropagationHelper.getUserFromUniqueID (uniqueid);

    // Retrieves the display name from the user registry based on the uniqueID.
    String securityName = reg.getUserSecurityName(uid);

    // Retrieves the groups associated with this uniqueID.
    java.util.List groupList = reg.getUniqueGroupIds(uid);

    // Creates the java.util.Hashtable with the information that you gathered
    // from the UserRegistry.
    java.util.Hashtable hashtable = new java.util.Hashtable();
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIALIAL_UNIQUEID, uniqueid);
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIALIAL_SECURITYNAME, securityName);
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIALIAL_GROUPS, groupList);

    // Adds a cache key that is used as part of the lookup mechanism for
    // the created Subject. The cache key can be an object, but has
    // an implemented toString method. Make sure the cacheKey contains enough
    // information to scope it to the user and any additional attributes you are
    // using. If you do not specify this property, the Subject is scoped to the
    // WSCREDENTIALIAL_UNIQUEID returned, by default.
    hashtable.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIALIAL_CACHE_KEY, "myCustomAttribute" + uniqueid);
    // Adds the hashtable to the shared state of the Subject.
    _sharedState.put(com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIALIAL_PROPERTIES_KEY, hashtable);
}
else if (requiresLogin == false)
{
    // For more information on this section, see

```

```

    // "Security attribute propagation" on page 448.
    // If you added a custom Token implementation, you can search through the
    // token holder list for it to deserialize.
    // Note: Any Java objects are automatically deserialized by
    // wsMapDefaultInboundLoginModule

for (int i=0; i<authzTokenList.size(); i++)
{
    TokenHolder tokenHolder = (TokenHolder) authzTokenList.get(i);
    if (tokenHolder.getName().equals("com.acme.MyCustomTokenImpl"))
    {
        byte[] myTokenBytes = tokenHolder.getBytes();

        // Passes these bytes into the constructor of your implementation
        // class for deserialization.
        com.acme.MyCustomTokenImpl myTokenImpl = new com.acme.MyCustomTokenImpl(myTokenBytes);
    }
}
}

public boolean commit() throws LoginException
{
    // (For more information on what to do during a commit, see
    // "Developing custom login modules for a system login configuration" on page 375.)
}

// Defines your login module variables
com.ibm.wsspi.security.token.AuthorizationToken customAuthzToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthToken = null;
java.util.Map _sharedState = null;
}

```

Configuring outbound mapping to a different target realm

By default, when WebSphere Application Server makes an outbound request from one server to another server in a different security realm, the request is rejected. This topic details alternatives for enabling one server to send outbound requests to a target server in a different realm.

About this task

This outbound request is rejected to protect against a rogue server reading potentially sensitive information if successfully impersonating the home of the object. Select one of the following alternative procedures so that one server can send outbound requests to a target server in a different realm. When you are finished with a procedure on the administrative console, click **Apply**.

- Do not perform mapping. Instead, allow the existing security information to flow to a trusted target server, even if the target server resides in a different realm. Complete the following steps in the administrative console:
 1. Click **Security > Global security**.
 2. Under RMI/IIOP security, click **CSlv2 outbound authentication**.
 3. Specify the target realms in the **Trusted target realms** field. You can specify each trusted target realm that is separated by a pipe (|) character. For example, specify *server_name.domain:port_number* for a Lightweight Directory Access Protocol (LDAP) server or the machine name for local operating system. If you want to propagate security attributes to a different target realm, you must specify that target realm in the **Trusted target realms** field.
- Use the Java Authentication and Authorization Service (JAAS) WSLogin application login configuration to create a basic authentication Subject that contains the credentials of the new target realm. This configuration enables you to log in with a realm, user ID, and password that are specific to the user registry of the target realm. You can provide the login information from within the Java Platform,

Enterprise Edition (Java EE) application that is making the outbound request or from within the RMI_OUTBOUND system login configuration. These two login options are described in the following information:

1. Use the WLogin application login configuration from within the Java EE application to log in and get a Subject that contains the user ID and the password of the target realm. The application can wrap the remote call with a WSubject.doAs call. For an example, see “Example: Using the WLogin configuration to create a basic authentication subject” on page 445.
2. Use the code sample in “Example: Using the WLogin configuration to create a basic authentication subject” on page 445 from this plug point within the RMI_OUTBOUND login configuration. Every outbound Remote Method Invocation (RMI) request passes through this login configuration when it is enabled. Complete the following steps to enable and plug in this login configuration:
 - a. Click **Security > Global security**.
 - b. Under RMI/IIOP security, click **CSiv2 outbound authentication**.
 - c. Select the **Custom outbound mapping** option. If the **Security Attribute Propagation** option is selected, then WebSphere Application Server is already using this login configuration and you do not need to enable custom outbound mapping.
 - d. Write a custom login module. For more information, see “Developing custom login modules for a system login configuration” on page 375.

The “Example: Sample login configuration for RMI_OUTBOUND” on page 446 shows a custom login module that determines whether the realm names match. In this example, the realm names do not match so the WLoginmodule is used to create a basic authentication Subject based on custom mapping rules. The custom mapping rules are specific to the customer environment and must be implemented using a realm to user ID and password mapping utility.
 - e. Configure the RMI_OUTBOUND login configuration so that your new custom login module is first in the list.
 - 1) Click **Security > Global security**.
 - 2) Under Java Authentication and Authorization Service, click **System logins > RMI_OUTBOUND**
 - 3) Under Additional Properties, click **JAAS login modules > New** to add your login module to the RMI_OUTBOUND configuration.
 - 4) Return to the JAAS login modules panel for RMI_OUTBOUND.
 - 5) Click **Set order** to change the order that the login modules are loaded so that your custom login is loaded first.

- Add the use_realm_callback and use_appcontext_callback options to the outbound mapping module for WLogin. To add these options, complete the following steps:
 1. Click **Security > Global security**.
 2. Under Java Authentication and Authorization Service, click **Application logins > WLogin**.
 3. Under Additional properties, click **JAAS login modules > com.ibm.ws.security.common.auth.module.WLoginModuleImpl**.
 4. Under Additional properties, click **Custom Properties > New**.
 5. On the Custom properties panel, enter use_realm_callback in the **Name** field and true in the **Value** field.
 6. Click **OK**.
 7. Click **New** to enter the second custom property.
 8. On the Custom properties panel, enter use_appcontext_callback in the **Name** field and true in the **Value** field.

The following changes are made to the security.xml file:

```
<entries xmi:id="JAASConfigurationEntry_2" alias="WLogin">
  <loginModules xmi:id="JAASLoginModule_2"
    moduleName="com.ibm.ws.security.common.auth.module.proxy.WLoginModuleProxy">
```

```

    authenticationStrategy="REQUIRED">
    <options xmi:id="Property_2" name="delegate"
      value="com.ibm.ws.security.common.auth.module.WSLoginModuleImpl"/>
    <options xmi:id="Property_3" name="use_realm_callback" value="true"/>
    <options xmi:id="Property_4" name="use_appcontext_callback" value="true"/>
  </loginModules>
</entries>

```

Example: Using the WSLogin configuration to create a basic authentication subject

This example shows how to use the WSLogin application login configuration from within a Java 2 Platform, Enterprise Edition (J2EE) application to log in and get a Subject that contains the user ID and the password of the target realm.

```

javax.security.auth.Subject subject = null;

try
{
    // Create a login context using the WSLogin login configuration and specify a
    // user ID, target realm, and password. Note: If the target_realm_name is the
    // same as the current realm, an authenticated Subject is created. However, if
    // the target_realm_name is different from the current realm, a basic
    // authentication Subject is created that is not validated. This unvalidated
    // Subject is created so that you can send a request to the different target
    // realm with valid security credentials for that realm.
    javax.security.auth.login.LoginContext ctx = new LoginContext("WSLogin",
        new WSCallbackHandlerImpl("userid", "target_realm_name", "password"));

    // Note: The following code is an alternative that validates the user ID and
    // password specified against the target realm. The code performs a remote call
    // to the target server and will return true if the user ID and password are
    // valid and false if the user ID and password are not valid. If false is
    // returned, a WSLoginFailedException exception is created. You can catch
    // that exception and perform a retry or stop the request from flowing by
    // allowing that exception to surface out of this login.

    // ALTERNATIVE LOGIN CONTEXT THAT VALIDATES THE USER ID AND PASSWORD TO THE
    // TARGET REALM

    /**** currently remarked out ****
    java.util.Map appContext = new java.util.HashMap();
        appContext.put(javax.naming.Context.INITIAL_CONTEXT_FACTORY,
            "com.ibm.websphere.naming.WsnInitialContextFactory");
        appContext.put(javax.naming.Context.PROVIDER_URL,
            "corbaloc:iiop:target_host:2809");

    javax.security.auth.login.LoginContext ctx = new LoginContext("WSLogin",
        new WSCallbackHandlerImpl("userid", "target_realm_name", "password", appContext));
    **** currently remarked out ****/

    // Starts the login
    ctx.login();

    // Gets the Subject from the context
    subject = ctx.getSubject();
}
catch (javax.security.auth.login.LoginException e)
{
    throw new com.ibm.websphere.security.auth.WSLoginFailedException (e.getMessage(), e);
}

if (subject != null)
{
    // Defines a privileged action that encapsulates your remote request.
    java.security.PrivilegedAction myAction = java.security.PrivilegedAction()
    {
        public Object run()

```



```

    {
        // Assumes a proxy is already defined. This example method returns a String
        return proxy.remoteRequest();
    }
});

// Starts this action using the basic authentication Subject needed for
// the target realm security requirements.
String myResult = (String) com.ibm.websphere.security.auth.WSSubject.doAs
    (subject, myAction);
}

```

Example: Sample login configuration for RMI_OUTBOUND

This example shows a sample login configuration for RMI_OUTBOUND that determines whether the realm names match between two servers.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        // (For more information on what to do during initialization, see
        // "Developing custom login modules for a system login configuration" on page 375.)
    }

    public boolean login() throws LoginException
    {
        // (For more information on what to do during login, see
        // "Developing custom login modules for a system login configuration" on page 375.)

        // Gets the WSPolicyCallback object
        Callback callbacks[] = new Callback[1];
        callbacks[0] = new com.ibm.wsspi.security.auth.callback.
            WSPolicyCallback("Protocol Policy Callback: ");

        try
        {
            callbackHandler.handle(callbacks);
        }
        catch (Exception e)
        {
            // Handles the exception
        }

        // Receives the RMI (CSiv2) policy object for checking the target realm
        // based upon information from the IOR.
        // Note: This object can be used to perform additional security checks.
        // See the application programming interface (API) documentation for
        // more information.
        csiv2PerformPolicy = (CSiv2PerformPolicy) ((WSPolicyCallback)callbacks[0]).
            getProtocolPolicy();

        // Checks if the realms do not match. If they do not match, then log in to
        // perform a mapping
        if (!csiv2PerformPolicy.getTargetSecurityName().equalsIgnoreCase(csiv2PerformPolicy.
            getCurrentSecurityName()))
        {
            try
            {
                // Do some custom realm -> user ID and password mapping
                MyBasicAuthDataObject myBasicAuthData = MyMappingLogin.lookup
                    (csiv2PerformPolicy.getTargetSecurityName());

                // Creates the login context with basic authentication data gathered from
                // custom mapping
                javax.security.auth.LoginContext ctx = new LoginContext("WSLogin",
                    new WSCallbackHandlerImpl(myBasicAuthData.userid,

```

```

        csiv2PerformPolicy.getTargetSecurityName(),
        myBasicAuthData.password));

// Starts the login
ctx.login();

        // Gets the Subject from the context. This subject is used to replace
        // the passed-in Subject during the commit phase.
        basic_auth_subject = ctx.getSubject();
    }
    catch (javax.security.auth.login.LoginException e)
    {
        throw new com.ibm.websphere.security.auth.
            WSLoginFailedException (e.getMessage(), e);
    }
}
}

public boolean commit() throws LoginException
{
    // (For more information on what to do during commit, see
    // "Developing custom login modules for a system login configuration" on page 375.)

    if (basic_auth_subject != null)
    {
        // Removes everything from the current Subject and adds everything from the
        // basic_auth_subject
        try
        {
            public final Subject basic_auth_subject_priv = basic_auth_subject;
            // Do this in a doPrivileged code block so that application code
            // does not need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.
                PrivilegedExceptionAction()
            {
                public Object run() throws WSLoginFailedException
                {
                    // Removes everything user-specific from the current outbound
                    // Subject. This a temporary Subject for this specific invocation
                    // so you are not affecting the Subject set on the thread. You may
                    // keep any custom objects that you want to propagate in the Subject.
                    // This example removes everything and adds just the new information
                    // back in.

                    try
                    {
                        subject.getPublicCredentials().clear();
                        subject.getPrivateCredentials().clear();
                        subject.getPrincipals().clear();
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }

                    // Adds everything from basic_auth_subject into the login subject.
                    // This completes the mapping to the new user.

                    try
                    {
                        subject.getPublicCredentials().addAll(basic_auth_subject.
                            getPublicCredentials());
                        subject.getPrivateCredentials().addAll(basic_auth_subject.
                            getPrivateCredentials());
                        subject.getPrincipals().addAll(basic_auth_subject.
                            getPrincipals());
                    }
                    catch (Exception e)
                    {

```

```

        throw new WLoginFailedException (e.getMessage(), e);
    }

    return null;
}
});
}
catch (PrivilegedActionException e)
{
    throw new WLoginFailedException (e.getException().getMessage(),
        e.getException());
}
}
}

// Defines your login module variables
com.ibm.wsspi.security.csiv2.CSiv2PerformPolicy csiv2PerformPolicy = null;
javax.security.auth.Subject basic_auth_subject = null;
}

```

Security attribute propagation

With *Security attribute propagation*, WebSphere Application Server can transport security attributes (authenticated Subject contents and security context information) from one server to another in your configuration. WebSphere Application Server might obtain these security attributes from either an enterprise user registry, which queries static attributes, or a custom login module, which can query static or dynamic attributes. Dynamic security attributes, which are custom in nature, might include the authentication strength that is used for the connection, the identity of the original caller, the location of the original caller, the IP address of the original caller, and so on.

Security attribute propagation provides propagation services using Java serialization for any objects that are contained in the Subject. However, Java code must be able to serialize and deserialize these objects. The Java programming language specifies the rules for how Java code can serialize an object. Because problems can occur when dealing with different platforms and versions of software, WebSphere Application Server also offers a token framework that enables custom serialization functionality. The token framework has other benefits that include the ability to identify the uniqueness of the token. This uniqueness determines how the Subject gets cached and the purpose of the token. The token framework defines four marker token interfaces that enable the WebSphere Application Server runtime to determine how to propagate the token.

Note: Any custom tokens that are used in this framework are not used by WebSphere Application Server for authorization or authentication. The framework serves as a way to notify WebSphere Application Server that you want these tokens propagated in a particular way. WebSphere Application Server handles the propagation details, but does not handle serialization or deserialization of custom tokens. Serialization and deserialization of these custom tokens are carried out by the implementation and handled by a custom login module.

With WebSphere Application Server Version 6.0 and later, a custom Java Authorization Contract for Container (JACC) provider can be configured to enforce access control for Java Platform, Enterprise Edition (Java EE) applications. A custom JACC provider can explore the custom security attributes in the caller JAAS subject in making access control decisions.

When a request is being authenticated, a determination is made by the login modules whether this request is an *initial login* or a *propagation login*. An initial login is the process of authenticating the user information, typically a user ID and password, and then calling the application programming interfaces (APIs) for the remote user registry to look up secure attributes that represent the user access rights. A propagation login is the process of validating the user information, typically a Lightweight Third Party Authentication (LTPA) token, and then deserializing a series of tokens that constitute both custom objects and token framework objects known to WebSphere Application Server.

The following marker tokens are introduced in the framework:

Authorization token

The authorization token contains most of the authorization-related security attributes that are propagated. The default authorization token is used by the WebSphere Application Server authorization engine to make Java Platform, Enterprise Edition (Java EE) authorization decisions. Service providers can use custom authorization token implementations to isolate their data in a different token, perform custom serialization and de-serialization, and make custom authorization decisions using the information in their token at the appropriate time. For information on how to use and implement this token type, see “Using the default propagation token” on page 459 and “Implementing a custom propagation token” on page 789.

Single sign-on (SSO) token

A custom SingleSignonToken token that is added to the Subject is automatically added to the response as an HTTP cookie and contains the attributes sent back to Web browsers. The token interface getName method with the getVersion method defines the cookie name. WebSphere Application Server defines a default SingleSignonToken token with the LtpaToken name and Version 2. The cookie name added is LtpaToken2. Do not add sensitive information, confidential information, or unencrypted data to the response cookie.

It is also recommended that any time that you use cookies, use the Secure Sockets Layer (SSL) protocol to protect the request. Using an SSO token, Web users can authenticate once when accessing Web resources across multiple WebSphere Application Servers. A custom SSO token extends this functionality by adding custom processing to the single sign-on scenario. For more information on SSO tokens, see “Implementing single sign-on to minimize Web user authentications” on page 297. For information on how to use and implement this token type, see “Using the default single sign-on token with default or custom token factory” on page 465 and “Implementing a custom single sign-on token” on page 806.

Propagation token

The propagation token is not associated with the authenticated user so it is not stored in the Subject. Instead, the propagation token is stored on the thread and follows the invocation wherever it goes. When a request is sent outbound to another server, the propagation tokens on that thread are sent with the request and the tokens are run by the target server. The attributes that are stored on the thread are propagated regardless of the Java Platform, Enterprise Edition (Java EE) RunAs user switches.

The default propagation token monitors and logs all user switches and host switches. You can add additional information to the default propagation token using the WSSecurityHelper application programming interfaces (APIs). To retrieve and set custom implementations of a propagation token, you can use the WSSecurityPropagationHelper class. For information on how to use and implement this token type, see “Using the default propagation token” on page 459 and “Implementing a custom propagation token” on page 789.

Authentication token

The authentication token flows to downstream servers and contains the identity of the user. This token type serves the same function as the Lightweight Third Party Authentication (LTPA) token in previous versions. Although this token type is typically reserved for internal WebSphere Application Server purposes, you can add this token to the Subject and the token is propagated using the getBytes method of the token interface.

A custom authentication token is used solely for the purpose of the service provider that adds it to the Subject. WebSphere Application Server does not use it for authentication purposes because a default authentication token exists that is used for WebSphere Application Server authentication. This token type is available for the service provider to identify how the custom data uses the token to perform custom authentication decisions. For information on how to use and implement this token type, see “Default authentication token” on page 452 and “Implementing a custom authentication token” on page 816.

KerberosTicket token

If Kerberos is the active Kerberos authentication mechanism, this class encapsulates a Kerberos ticket and associated information as viewed from the point of view of the client.

Horizontal propagation versus downstream propagation

In WebSphere Application Server, both horizontal propagation, which uses single sign-on for Web requests, and downstream propagation, which uses Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) to access enterprise beans, are available.

Horizontal propagation

In horizontal propagation, security attributes are propagated among front-end servers. The serialized security attributes, which are the Subject contents and the propagation tokens, can contain both static and dynamic attributes. The single sign-on (SSO) token stores additional system-specific information that is needed for horizontal propagation. The information contained in the SSO token tells the receiving server where the originating server is located and how to communicate with that server. Additionally, the SSO token also contains the key to look up the serialized attributes. To enable horizontal propagation, you must configure the single sign-on token and the Web inbound security attribute propagation features. You can configure both of these features using the administrative console.

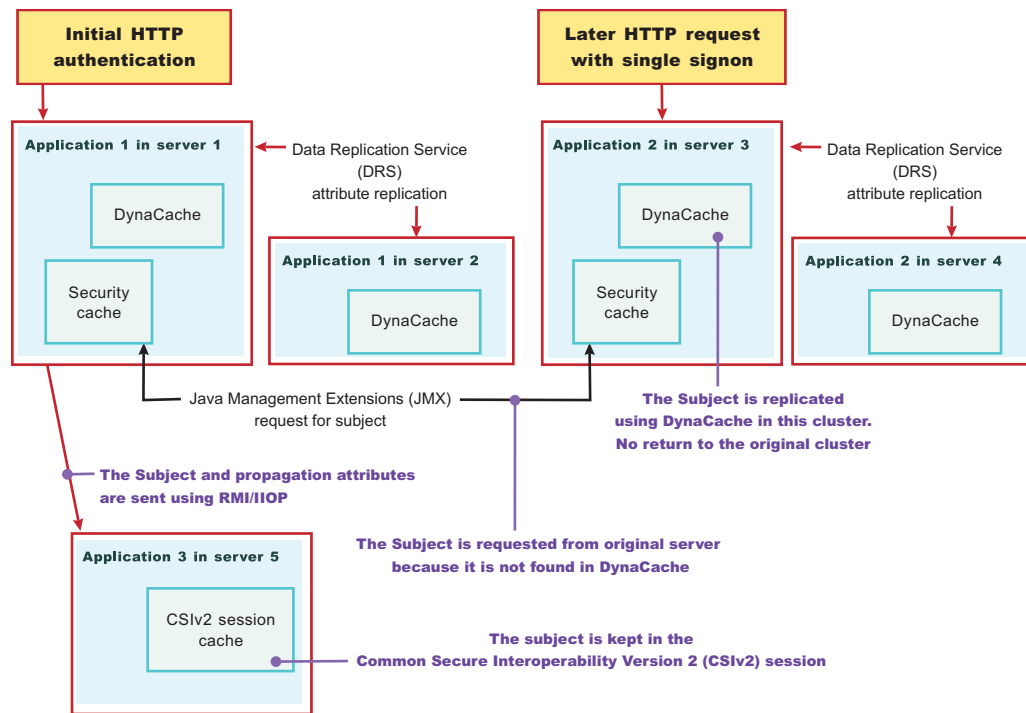
When front-end servers are configured and in the same data replication service (DRS) replication domain, the application server automatically propagates the serialized information to all of the servers within the same domain. In figure 1, application 1 is deployed on server 1 and server 2, and both servers are members of the same DRS replication domain. If a request originates from application 1 on server 1 and then gets redirected to application 1 on server 2, the original login attributes are found on server 2 without additional remote requests.

However, if the request originates from application 1 on either server 1 or server 2, but the request is redirected to application 2 on either server 1 or server 2, the serialized information is not found in the DRS cache because the servers are not configured in the same replication domain. As a result, a remote Java Management Extensions (JMX) request is sent back to the originating server that hosts application 1 to obtain the serialized information so that original login information is available to the application. By getting the serialized information using a single JMX remote call back to the originating server, the following benefits are realized:

- You gain the function of retrieving login information from the original server.
- You do not need to perform any remote user registry calls because the application server can regenerate the Subject from the serialized information. Without this ability, the application server might make five to six separate remote calls.

Figure 1

1. User authenticates to server 1.
2. Server 1 makes an RMI request to server 5.
3. User accesses another Web application on server 3.



Performance implications for horizontal propagation

The performance implications of either the DRS or JMX remote call depends upon your environment. The DRS or JMX remote call is used for obtaining the original login attributes. Horizontal propagation reduces many of the remote user registry calls in cases where these calls cause the most performance problems for an application. However, the de-serialization of these objects also might cause performance degradation, but this degradation might be less than the remote user registry calls. It is recommended that you test your environment with horizontal propagation enabled and disabled. In cases where you must use horizontal propagation for preserving original login attributes, test whether DRS or JMX provides better performance in your environment. Typically, it is recommended that you configure DRS both for failover and performance reasons. However, because DRS propagates the information to all of the servers in the same replication domain (whether the servers are accessed or not), there might be a performance degradation if too many servers are in the same replication domain. In this case, either reduce the number of servers in the replication domain or do not configure the servers in a DRS replication domain. The later suggestion causes a JMX remote call to retrieve the attributes, when needed, which might be quicker overall.

Downstream propagation

In *downstream propagation*, a Subject is generated at the Web front-end server, either by a propagation login or a user registry login. WebSphere Application Server propagates the security information downstream for enterprise bean invocations when both Remote Method Invocation (RMI) outbound and inbound propagation are enabled.

Benefits of propagating security attributes

The security attribute propagation feature of WebSphere Application Server has the following benefits:

- Enables WebSphere Application Server to use the security attribute information for authentication and authorization purposes. The propagation of security attributes can eliminate the need for user registry

calls at each remote hop along an invocation. Previous versions of WebSphere Application Server propagated only the user name of the authenticated user, but ignored other security attribute information that needed to be regenerated downstream using remote user registry calls. To accentuate the benefits of this new functionality, consider the following example:

In previous releases, you might use a reverse proxy server (RPSS), such as WebSEAL, to authenticate the user, gather group information, and gather other security attributes. As stated previously, WebSphere Application Server accepted the identity of the authenticated user, but disregarded the additional security attribute information. To create a Java Authentication and Authorization Service (JAAS) Subject containing the needed WSCredential and WSPincipal objects, WebSphere Application Server made 5 to 6 calls to the user registry. The WSCredential object contains various security information that is required to authorize a Java EE resource. The WSPincipal object contains the realm name and the user that represents the principal for the Subject.

In the current release of the Application Server, information that is obtained from the reverse proxy server can be used by WebSphere Application Server and propagated downstream to other server resources without additional calls to the user registry. The retaining of the security attribute information enables you to protect server resources properly by making appropriate authorization and trust-based decisions. User switches that occur because of Java EE RunAs configurations do not cause the application server to lose the original caller information. This information is stored in the PropagationToken located on the running thread.

- Enables third-party providers to plug in custom tokens. The token interface contains a getBytes method that enables the token implementation to define custom serialization, encryption methods, or both.
- Provides the ability to have multiple tokens of the same type within a Subject created by different providers. WebSphere Application Server can handle multiple tokens for the same purpose. For example, you might have multiple authorization tokens in the Subject and each token might have distinct authorization attributes that are generated by different providers.
- Provides the ability to have a unique ID for each token type that is used to formulate a more unique subject identifier than just the user name in cases where dynamic attributes might change the context of a user login. The token type has a getUniqueId() method that is used for returning a unique string for caching purposes. For example, you might need to propagate a location ID, which indicates the location from which the user logs into the system. This location ID can be generated during the original login using either an reverse proxy server or the WEB_INBOUND login configuration and added to the Subject prior to serialization. Other attributes might be added to the Subject as well and use a unique ID. All of the unique IDs must be considered for the uniqueness of the entire Subject. WebSphere Application Server has the ability to specify what is unique about the information in the Subject, which might affect how the user accesses the Subject later.

Default authentication token

Do not use the default authentication token in service provider code. This default token is used by the WebSphere Application Server run-time code only and is authentication mechanism specific.

Any modifications to this token by service provider code can potentially cause interoperability problems. If you need to create an authentication token for custom usage, see “Implementing a custom authentication token” on page 816 for more information.

Changing the token factory that is associated with the default authentication token

When WebSphere Application Server generates a default authentication token, the application server utilizes the TokenFactory class that is specified using the com.ibm.wsspi.security.token.authenticationTokenFactory property. To modify this property using the administrative console, complete the following steps:

1. Click **Security > Global security**.
2. Under Additional properties, click **Custom properties**.

The `com.ibm.ws.security.ltpa.LTPATokenFactory` token factory is the default for this property. The `LTPATokenFactory` token factory uses the `DESede/ECB/PKCS5Padding` cipher. This token factory creates an interoperable Lightweight Third Party Authentication (LTPA) token. If you change this token factory, you lose the interoperability with any servers running a version of WebSphere Application Server prior to Version 5.1.1 and any other servers that do not support the new token factory implementation. However, if all of your application servers use WebSphere Application Server Version 5.1.1 or later and all of your servers use your new token factory, this interoperability is not a problem.

If you associate the `com.ibm.ws.security.ltpa.LTPAToken2Factory` token factory with the `com.ibm.wsspi.security.token.authenticationTokenFactory` property, the token is Advanced Encryption Standard (AES) encrypted. However, you need to weigh the performance against your security needs. You might add additional attributes to the authentication token in the Subject during a login that are available downstream.

If you need to perform your own signing and encryption of the default authentication token, you must implement the following classes:

- `com.ibm.wsspi.security.ltpa.Token`
- `com.ibm.wsspi.security.ltpa.TokenFactory`

Your token factory implementation instantiates (`createToken`) and validates (`validateTokenBytes`) your token implementation. You can use the LTPA keys that are passed into the `initialize` method of the token factory or you can use your own keys. If you use your own keys, they must be the same everywhere to validate the tokens that are generated using those keys. See the API documentation, available through a link on the front page of the information center, for more information on implementing your own custom token factory. To associate your token factory with the default authentication token using the administrative console, complete the following steps:

1. Click **Security > Global security**.
2. Under Additional properties, click **Custom properties**.
3. Locate the `com.ibm.wsspi.security.token.authenticationTokenFactory` property and verify that the value of this property matches your custom token factory implementation.
4. Verify that your implementation classes are put into the `install_dir/classes` directory so that the WebSphere Application Server class loader can load the classes.

Propagating security attributes among application servers

Use the security attribute propagation feature of WebSphere Application Server to send security attribute information regarding the original login to other servers using a token. This topic will help to configure WebSphere Application Server to propagate security attributes to other servers.

About this task

To fully enable security attribute propagation, you must configure the single sign-on (SSO), Common Secure Interoperability Version 2 (CSIv2) inbound, and CSIv2 outbound panels in the WebSphere Application Server administrative console. You can enable just the portions of security attribute propagation relevant to your configuration. For example, you can enable Web propagation, which is propagation amongst front-end application servers, using either the push technique (DynaCache) or the pull technique (remote method to originating server).

You also can choose whether to enable Remote Method Invocation (RMI) outbound and inbound propagation, which is commonly called downstream propagation. Typically both types of propagation are enabled for any given cell. In some cases, you might want to choose a different option for a specific application server using the server security panel within the specific application server settings.

Note: To prevent propagating the same security attributes among application servers multiple times, WebSphere Application Server verifies that a Lightweight Third Party Authentication (LTPA) token

does not exist. Two cases can occur. Absence of the LTPA token tells the Application Server that propagation can proceed. Presence of the LTPA token indicates that propagation has occurred if the LTPA token has been generated within the cluster. However, in the second case, if the LTPA token is present, but has been generated by a server outside the cluster, such as by Tivoli Access Manager, Lotus Domino or a different Application Server cluster, security attributes are not propagated.

To access the server security panel in the administrative console, click **Servers > Application Servers > *server_name***. Under Security, click **Server security**.

Complete the following steps to configure WebSphere Application Server for security attribute propagation:

1. Access the WebSphere Application Server administrative console by typing `http://server_name:port_number/ibm/console`. The administrative console address might differ if you have previously changed the port number.
2. Click **Security > Global security**.
3. Under Web security, click **Single sign-on (SSO)**.
4. Optional: Select the **Interoperability Mode** option if you need to interoperate with servers that do not support security attribute propagation. Servers that do not support security attribute propagation receive the Lightweight Third Party Authentication (LTPA) token and the Propagation token, but ignore the security attribute information that they do not understand.
5. Select the **Web inbound security attribute propagation** option. The Web inbound security attribute propagation option enables horizontal propagation, which allows the receiving SSO token to retrieve the login information from the original login server. If you do not enable this option, downstream propagation can occur if you enable the Security Attribute Propagation option on both the CSiv2 Inbound authentication and CSiv2 outbound authentication panels.

Typically, you enable the Web inbound security attribute propagation option if you need to gather dynamic security attributes set at the original login server that cannot be regenerated at the new front-end server. These attributes include any custom attributes that might be set in the PropagationToken token using the `com.ibm.websphere.security.WSSecurityHelper` application programming interfaces (APIs). You must determine whether enabling this option improves or degrades the performance of your system. While the option prevents some remote user registry calls, the deserialization and decryption of some tokens might impact performance. In some cases propagation is faster, especially if your user registry is the bottleneck of your topology. It is recommended that you measure the performance of your environment both using and not using this option. When you test the performance, it is recommended that you test in the operating environment of the typical production environment with the typical number of unique users accessing the system simultaneously.

6. Click **Security > Global security**. Under RMI/IIOP security, click **CSiv2 inbound authentication**. The Login configuration field specifies `RMI_INBOUND` as the system login configuration that is used for inbound requests. To add custom Java Authentication and Authorization Service (JAAS) login modules, complete the following steps:
 - a. Click **Security > Global security**. Under Java Authentication and Authorization Service, click **System logins**. A list of the system login configurations is displayed. WebSphere Application Server provides the following pre-configured system login configurations: `DEFAULT`, `LTPA`, `LTPA_WEB`, `RMI_INBOUND`, `RMI_OUTBOUND`, `SWAM`, `WEB_INBOUND`, `wssecurity.IDAssertion`, and `wssecurity.Signature`. Do not delete these predefined configurations.

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

- b. Click the name of the login configuration that you want to modify.
- c. Under Additional Properties, click **JAAS Login Modules**. The JAAS Login Modules panel is displayed, which lists all of the login modules that are processed in the login configuration. Do not

delete the required JAAS login modules. Instead, you can add custom login modules before or after the required login modules. If you add custom login modules, do not begin their names with `com.ibm.ws.security.server`.

You can specify the order in which the login modules are processed by clicking **Set Order**.

7. Select the **Security attribute propagation** option on the CSiv2 inbound authentication panel. When you select **Security Attribute Propagation**, the server advertises to other application servers that it can receive propagated security attributes from another server in the same realm over the Common Secure Interoperability version 2 (CSiv2) protocol.
8. Click **Security > Global security**. Under RMI/IIOP security, click **CSiv2 Outbound authentication**. The CSiv2 outbound authentication panel is displayed. The **Login configuration** field specifies `RMI_OUTBOUND` as the JAAS login configuration that is used for outbound configuration. You cannot change this login configuration. Instead, you can customize this login configuration by completing the substeps that are listed previously for CSiv2 Inbound authentication.
9. Optional: Verify that the **Security Attribute Propagation** option is selected if you want to enable outbound Subject and security context token propagation for the Remote Method Invocation (RMI) protocol. When you select this option, WebSphere Application Server serializes the Subject contents and the PropagationToken contents. After the contents are serialized, the server uses the CSiv2 protocol to send the Subject and PropagationToken token to the target servers that support security attribute propagation. If the receiving server does not support security attribute tokens, WebSphere Application Server sends the Lightweight Third Party Authentication (LTPA) token only.

Note: WebSphere Application Server propagates only the objects within the Subject that it can serialize. The server propagates custom objects on a best-effort basis.

When **Security Attribute Propagation** is enabled, WebSphere Application Server adds marker tokens to the Subject to enable the target server to add additional attributes during the inbound login. During the commit phase of the login, the marker tokens and the Subject are marked as read-only and cannot be modified thereafter.

10. Optional: Select the **Custom Outbound Mapping** option if you clear the **Security Attribute Propagation** option and you want to use the `RMI_OUTBOUND` login configuration. If neither the **Custom Outbound Mapping** option nor the **Security Attribute Propagation** option is selected, WebSphere Application Server does not call the `RMI_OUTBOUND` login configuration. If you need to plug in a credential mapping login module, you must select the **Custom Outbound Mapping** option.
11. Optional: Specify trusted target realm names in the **Trusted Target Realms** field. By specifying these realm names, information can be sent to servers that reside outside the realm of the sending server to support inbound mapping that is at these downstream servers. To perform outbound mapping to a realm different from the current realm, you must specify the realm in this field so that you can get to this point without having the request rejected because of a realm mismatch. If you need WebSphere Application Server to propagate security attributes to another realm when a request is sent, you must specify the realm name in the **Trusted Target Realms** field. Otherwise, the security attributes are not propagated to the unspecified realm. You can add multiple target realms by adding a pipe (|) delimiter between each entry.
12. Optional: Enable propagation for a pure client. For a pure client to propagate attributes added to the invocation Subject, you must add the following property to the `sas.client.props` file:
`com.ibm.CSI.rmiOutboundPropagationEnabled=true`

Note: The `sas.client.props` file is located at `<WAS-HOME>/profiles/<ProfileName>/properties`.

Results

After completing these steps, you have configured WebSphere Application Server to propagate security attributes to other servers.

What to do next

If you need to disable security attribute propagation, determine whether you need to disable it for either the server level or the cell level.

Note: Changes to the server-level settings override the cell settings.

To disable security attribute propagation on the server level, complete the following steps:

1. Click **Server > Application Servers > *server_name***.
2. Under Security, click **Server security**.
3. Select the **RMI/IIOP security for this server overrides cell settings** option.
4. Disable security attribute propagation for inbound requests by clicking **CSI inbound authentication** under Additional Properties and clearing the **Security attribute propagation** option.
5. Disable security attribute propagation for outbound requests by clicking **CSI outbound authentication** under Additional Properties and clearing the **Security attribute propagation** option.

To disable security attribute propagation on the cell level, undo each of the steps that you completed to enable security attribute propagation in this task.

Using the default authorization token

This topic explains how WebSphere Application Server uses the default authorization token. Consider using the default authorization token when you are looking for a place to add string attributes that get propagated downstream.

About this task

However, make sure that the attributes you add to the authorization token are specific to the user that is associated with the authenticated Subject. If they are not specific to a user, the attributes probably belong in the propagation token, which is also propagated with the request. For more information on the propagation token, see “Using the default propagation token” on page 459. To add attributes into the authorization token, you must plug in a custom login module into the various system login modules that are configured. Any login module configuration that has the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` implementation configured can receive propagated information and can generate propagation information that can be sent outbound to another server.

If propagated attributes are not presented to the login configuration during an initial login, a default authorization token is created in the `wsMapDefaultInboundLoginModule` login module after the login occurs in the `ltpaLoginModule` login module. You can obtain a reference to the default authorization token from the login method using the `sharedState` hashmap. You must plug in the custom login module after the `wsMapDefaultInboundLoginModule` implementation for WebSphere Application Server to see the default authorization token.

For more information on the Java Authentication and Authorization Service (JAAS) programming model, see *Security: Resources for learning*.

- Obtain a reference to the default authorization token from the login method.
- Add attributes to the token.
- Read existing attributes used for authorization.
- Ensure that your custom login module code is trusted. Whenever you plug a custom login module into the WebSphere Application Server login infrastructure, you must ensure that the code is trusted. When you add the login module into the `app_server_root/classes` directory, it has Java 2 Security `AllPermissions` permissions. It is recommended that you add your login module and other infrastructure classes into a private directory. However, if you use a private directory, modify the

\$(WAS_INSTALL_ROOT)/properties/server.policy file so that the private directory, Java archive (JAR) file, or both have the permissions that are needed to run the application programming interfaces (API) that are called from the login module. Because the login module might run after the application code on the call stack, you might consider adding a doPrivileged code block so that you do not need to add additional permissions to your applications.

- Modify the authorization token factory to use a token factory other than the default token factory.

When WebSphere Application Server generates a default authorization token, the application server utilizes the TokenFactory class that is specified using the com.ibm.wsspi.security.token.authorizationTokenFactory property.

The com.ibm.ws.security.ltpa.AuthzPropTokenFactory token factory is the default. This token factory encodes the data, but does not encrypt the data in the authorization token. Because the authorization token typically flows over Common Secure Interoperability Version 2 (CSIv2) using Secure Sockets Layer (SSL), encrypting the token is not necessary. However, if you need additional security for the authorization token, you can associate a different token factory implementation with this property to get encryption. For example, if you associate the com.ibm.ws.security.ltpa.LTPAToken2Factory token factory with this property, the token uses Advanced Encryption Standard (AES) encryption. However, you need to weigh the performance impacts against your security needs. Adding sensitive information to the authorization token is one reason to change the token factory implementation to something that encrypts rather than just encodes.

1. Open the administrative console.
2. Click **Security > Global security**.
3. Under Additional properties, click **Custom properties**.

- Perform your own signing and encryption of the default authorization token.

If you want to perform your own signing and encryption of the default authorization token, you must implement the following classes:

- com.ibm.wsspi.security.ltpa.Token
- com.ibm.wsspi.security.ltpa.TokenFactory

Your token factory implementation instantiates and validates your token implementation. You can use the Lightweight Third Party Authentication (LTPA) keys that are passed into the initialize method of the token factory or you can use your own keys. If you use your own keys, they must be the same everywhere to validate the tokens that are generated using those keys. See the API documentation, that is available through a link on the front page of the information center, for more information on implementing your own custom token factory.

- Associate your token factory with the default authorization token.

To associate your token factory with the default authorization token, using the administrative console, complete the following steps:

1. Click **Security > Global security**.
2. Under Additional properties, click **Custom properties**.
3. Locate the com.ibm.wsspi.security.token.authorizationTokenFactory property and verify that the value of this property matches your custom token factory implementation.
4. Verify that your implementation classes are put into the *app_server_root/classes* directory so that the WebSphere Application Server class loader can load the classes.
5. Verify that your implementation classes are put into the \${USER_INSTALL_ROOT}/classes directory so that the WebSphere Application Server class loader can load the classes.

Example

The following example shows the complete task of obtaining a reference to the default authorization token from the login method, adding attributes to the token, and reading from the existing attributes that are used for authorization.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        // (For more information on initialization, see
        // "Developing custom login modules for a system login configuration" on page 375.)

        // Get a reference to the sharedState map that is passed in during initialization.
        _sharedState = sharedState;
    }

    public boolean login() throws LoginException
    {
        // (For more information on what to do during login, see
        // "Developing custom login modules for a system login configuration" on page 375.)

        // Look for the default AuthorizationToken in the shared state
        defaultAuthzToken = (com.ibm.wsspi.security.token.AuthorizationToken)
            sharedState.get
            (com.ibm.wsspi.security.auth.callback.Constants.WSAUTHZTOKEN_KEY);

        // Might not always have one of these generated. It depends on the login
        // configuration setup.
        if (defaultAuthzToken != null)
        {
            try
            {
                // Add a custom attribute
                defaultAuthzToken.addAttribute("key1", "value1");

                // Determine all of the attributes and values that exist in the token.
                java.util.Enumeration listOfAttributes = defaultAuthzToken.
                    getAttributeNames();

                while (listOfAttributes.hasMoreElements())
                {
                    String key = (String) listOfAttributes.nextElement();

                    String[] values = (String[]) defaultAuthzToken.getAttributes (key);

                    for (int i=0; i<values.length; i++)
                    {
                        System.out.println ("Key: " + key + ", Value[" + i + "]: "
                            + values[i]);
                    }
                }

                // Read the existing uniqueID attribute.
                String[] uniqueID = defaultAuthzToken.getAttributes
                    (com.ibm.wsspi.security.token.AttributeNameConstants.
                        WSCREDENTIAL_UNIQUEID);

                // Get the uniqueID from the String[]
                String unique_id = (uniqueID != null &&
                    uniqueID[0] != null) ? uniqueID[0] : "";

                // Read the existing expiration attribute.
                String[] expiration = defaultAuthzToken.getAttributes
                    (com.ibm.wsspi.security.token.AttributeNameConstants.
                        WSCREDENTIAL_EXPIRATION);

                // An example of getting a long expiration value from the string array.
                long expire_time = 0;
                if (expiration != null && expiration[0] != null)
                    expire_time = Long.parseLong(expiration[0]);
            }
        }
    }
}

```



```

// Read the existing display name attribute.
String[] securityName = defaultAuthzToken.getAttributes
    (com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_SECURITYNAME);

// Get the display name from the String[]
String display_name = (securityName != null &&
    securityName[0] != null) ? securityName[0] : "";

// Read the existing long securityName attribute.
String[] longSecurityName = defaultAuthzToken.getAttributes
    (com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_LONGSECURITYNAME);

// Get the long security name from the String[]
String long_security_name = (longSecurityName != null &&
    longSecurityName[0] != null) ? longSecurityName[0] : "";

// Read the existing group attribute.
String[] groupList = defaultAuthzToken.getAttributes
    (com.ibm.wsspi.security.token.AttributeNameConstants.
        WSCREDENTIAL_GROUPS);

// Get the groups from the String[]
ArrayList groups = new ArrayList();
if (groupList != null)
{
    for (int i=0; i<groupList.length; i++)
    {
        System.out.println ("group[" + i + "] = " + groupList[i]);
        groups.add(groupList[i]);
    }
}
catch (Exception e)
{
    throw new WSLoginFailedException (e.getMessage(), e);
}
}

public boolean commit() throws LoginException
{
    // (For more information on what to do during commit, see
    // "Developing custom login modules for a system login configuration" on page 375.)
}

private java.util.Map _sharedState = null;
private com.ibm.wsspi.security.token.AuthorizationToken defaultAuthzToken = null;
}

```

Using the default propagation token

A default propagation token is located on the running thread for applications and the security infrastructure to use. The product propagates this default propagation token downstream and the token stays on the thread where the invocation lands at each hop.

About this task

The data is available from within the container of any resource where the propagation token lands. Remember that you must enable the propagation feature at each server where a request is sent for propagation to work. Make sure that you enable security attribute propagation for all of the cells in your environment where you want propagation.

There is a `WSSecurityHelper` class that has application programming interfaces (APIs) for accessing the `PropagationToken` attributes. This topic documents the usage scenarios and includes examples. A close relationship exists between the propagation token and the work area feature. The main difference between these features is that after you add attributes to the propagation token, you cannot change the attributes. You cannot change these attributes so that the security runtime can add auditable information and have that information remain there for the life of the invocation. Any time that you add an attribute to a specific key, an `ArrayList` object is stored to hold that attribute. Any new attribute that is added with the same key is added to the `ArrayList` object. When you call `getAttributes`, the `ArrayList` object is converted to a `String` array and the order is preserved. The first element in the `String` array is the first attribute added for that specific key.

In the default propagation token, a change flag is kept that logs any data changes to the token. These changes are tracked to enable WebSphere Application Server to know when to send the authentication information downstream again so that the downstream server has those changes. Normally, Common Secure Interoperability Version 2 (CSIv2) maintains a session between servers for an authenticated client. If the propagation token changes, a new session is generated and subsequently a new authentication occurs. Frequent changes to the propagation token during a method cause frequent downstream calls. If you change the token prior to making many downstream calls or you change the token between each downstream call, you might impact security performance.

- Obtain the server list from the default propagation token.

Every time the propagation token is propagated and used to create the authenticated Subject, either horizontally or downstream, the name of the receiving application server is logged into the propagation token. The format of the host is "Cell:Node:Server", which provides you access to the cell name, node name, and server name of each application server that receives the invocation.

The following code provides you with this list of names and can be called from a Java 2 Platform, Enterprise Edition (J2EE) application.

The format of each server in the list is: *cell:node_name:server_name*. The output, for example, is: `myManager:node1:server1`

```
String[] server_list = null;

// If security is disabled on this application server, do not bother checking
// if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        // Gets the server_list string array
        server_list = com.ibm.websphere.security.WSSecurityHelper.getServerList();
    }
    catch (Exception e)
    {
        // Performs normal exception handling for your application
    }

    if (server_list != null)
    {
        // print out each server in the list, server_list[0] is the first server
        for (int i=0; i<server_list.length; i++)
```

```

    {
        System.out.println("Server[" + i + "] = " + server_list[i]);
    }
}
}

```

- Obtain the list of callers, using the `getCallerList` API.

A default propagation token is generated any time an authenticated user is set on the running thread or anyone tries to add attributes to the propagation token. Whenever an authenticated user is set on the thread, the user is logged in the default propagation token. At times, the same user might be logged in multiple times if the `RunAs` user is different from the caller. The following list provides the rules that are used to determine if a user that is added to the thread gets logged into the propagation token:

- The current Subject must be authenticated. For example, an unauthenticated Subject is not logged.
- The current authenticated Subject is logged if a Subject is not previously logged.
- The current authenticated Subject is logged if the last authenticated Subject that is logged does not contain the same user.
- The current authenticated Subject is logged on each unique application server that is involved in the propagation process.

The following code sample shows how to use the `getCallerList` API.

The format of each caller in the list is: *cell:node_name:server_name:realm:port_number/securityName*. The output, for example, is: `myManager:node1:server1:ldap.austin.ibm.com:389/jsmith`.

```

String[] caller_list = null;

// If security is disabled on this application server, do not check the caller list
if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        // Gets the caller_list string array
        caller_list = com.ibm.websphere.security.WSSecurityHelper.getCallerList();
    }
    catch (Exception e)
    {
        // Performs normal exception handling for your application
    }

    if (caller_list != null)
    {
        // Prints out each caller in the list, caller_list[0] is the first caller
        for (int i=0; i<caller_list.length;i++)
        {
            System.out.println("Caller[" + i + "] = " + caller_list[i]);
        }
    }
}

```

- Obtain the security name of the first authenticated user, using the `getFirst Caller` API.

Whenever you want to know which authenticated caller started the request, you can call the `getFirstCaller` method and the caller list is parsed. However, this method returns the security name of the caller only. If you need to know more than the security name, call the `getCallerList` method and retrieve the first entry in the String array. This entry provides all the caller information.

The following code sample retrieves the security name of the first authenticated caller using the `getFirstCaller` API.

The output, for example, is: `jsmith`.

```
String first_caller = null;

// If security is disabled on this application server, do not bother checking
// if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        // Gets the first caller
        first_caller = com.ibm.websphere.security.WSSecurityHelper.getFirstCaller();

        // Prints out the caller name
        System.out.println("First caller: " + first_caller);
    }
    catch (Exception e)
    {
        // Performs normal exception handling for your application
    }
}
}
```

- Obtain the name of the first application server for a request, using the `getFirstServer` method. Whenever you want to know what the first application server is for this request, call the `getFirstServer` method directly.

The following code sample retrieves the name of the first application server using the `getFirstServer` API.

The output, for example, is: `myManager:node1:server1`.

```
String first_server = null;

// If security is disabled on this application server, do not bother checking
// if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        // Gets the first server
        first_server = com.ibm.websphere.security.WSSecurityHelper.getFirstServer();

        // Prints out the server name
        System.out.println("First server: " + first_server);
    }
    catch (Exception e)
    {
        // Performs normal exception handling for your application
    }
}
}
```

- Add custom attributes to the default propagation token, using the `addPropagationAttribute` API. You can add custom attributes to the default propagation token for application usage. This token follows the request downstream so that the attributes are available when needed. When you use the default propagation token to add attributes, you must understand the following issues:
 - Adding information to the propagation token affects CSiv2 session caching. Add information sparingly between remote requests.
 - After you add information with a specific key, the information cannot be removed.
 - You can add as many values to a specific key as you need. However, all of the values must be available from a returned String array in the order that they were added.
 - The propagation token is available only on servers where propagation and security are enabled.

- The Java 2 Security `javax.security.auth.AuthPermission wssecurity.addPropagationAttribute` attribute is needed to add attributes to the default propagation token.
- An application cannot use keys that begin with either `com.ibm.websphere.security` or `com.ibm.wsspi.security`. These prefixes are reserved for system usage.

The following code sample shows how to use the `addPropagationAttribute` API. See “Using the default propagation token” on page 459 to retrieve attributes using the `getPropagationAttributes` application programming interface (API).

```
// If security is disabled on this application server,
// do not check the status of server security
if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        // Specifies the key and values
        String key = "mykey";
        String value1 = "value1";
        String value2 = "value2";

        // Sets key, value1
        com.ibm.websphere.security.WSSecurityHelper.
            addPropagationAttribute (key, value1);

        // Sets key, value2
        String[] previous_values = com.ibm.websphere.security.WSSecurityHelper.
            addPropagationAttribute (key, value2);

        // Note: previous_values should contain value1
    }
    catch (Exception e)
    {
        // Performs normal exception handling for your application
    }
}
```

- Obtain your custom attributes with the `getPropagationAttributes` API.

Custom attributes are added to the default propagation token using the `addPropagationAttribute` API. Retrieve these attributes using the `getPropagationAttributes` API. This token follows the request downstream so the attributes are available when needed. When you use the default propagation token to retrieve attributes, you must understand the following issues:

- The propagation token is available only on servers where propagation and security are enabled.
- The Java 2 Security `javax.security.auth.AuthPermission "wssecurity.getPropagationAttributes"` permission is needed to retrieve attributes from the default propagation token.

See [Adding custom attributes to the default PropagationToken](#) to add attributes using the `addPropagationAttributes` API.

The following code sample shows how to use the `getPropagationAttributes` API.

```
// If security is disabled on this application server, do not bother checking
if (com.ibm.websphere.security.WSSecurityHelper.isServerSecurityEnabled())
{
    try
    {
        String key = "mykey";
        String[] values = null;
    }
}
```

```

// Sets key, value1
    values = com.ibm.websphere.security.WSSecurityHelper.
        getPropagationAttributes (key);

// Prints the values
for (int i=0; i<values.length; i++)
{
    System.out.println("Value[" + i + "] = " + values[i]);
}
}
catch (Exception e)
{
    // Performs normal exception handling for your application
}
}

```

The output, for example, is:

```

Value[0] = value1
Value[1] = value2

```

- Modify the propagation token factory configuration to use a token factory other than the default token factory.

When WebSphere Application Server generates a default propagation token, the Application Server utilizes the TokenFactory class that is specified using the `com.ibm.wsspi.security.token.propagationTokenFactory` property.

The default token factory that is specified for this property is called `com.ibm.ws.security.ltpa.AuthzPropTokenFactory`. This token factory encodes the data in the propagation token and does not encrypt the data. Because the propagation token typically flows over CSiv2 using Secure Sockets Layer (SSL), encrypting the token is not required. However, if you need additional security for the propagation token, you can associate a different token factory implementation with this property to get encryption. For example, if you choose to associate the `com.ibm.ws.security.ltpa.LTPAToken2Factory` token factory with this property, the token is AES encrypted. However, you need to weigh the performance impacts against your security needs. Adding sensitive information to the propagation token is a good reason to change the token factory implementation to something that encrypts rather than just encodes.

1. Open the administrative console.
 2. Click **Security > Global security**.
 3. Click **Custom properties**.
- Perform your own signing and encryption of the default propagation token.
If you want to perform your own signing and encryption of the default propagation token, you must implement the following classes:
 - `com.ibm.wsspi.security.ltpa.Token`
 - `com.ibm.wsspi.security.ltpa.TokenFactory`

Your token factory implementation instantiates and validates your token implementation. You can choose to use the Lightweight Third Party Authentication (LTPA) keys and have them pass into the `initialize` method of the token factory, or you can use your own keys. If you use your own keys, they must be the same everywhere to validate the tokens that are generated using those keys. See the API documentation, available through a link on the front page of the information center, for more information on implementing your own custom token factory.

- Associate your token factory with the default propagation token.
 1. Open the administrative console.
 2. Click **Security > Global security**.
 3. Click **Custom properties**.

4. Locate the `com.ibm.wsspi.security.token.propagationTokenFactory` property and verify that the value of this property matches your custom token factory implementation.
5. Verify that your implementation classes are put into the `app_server_root/classes` directory so that the WebSphere Application Server class loader can load the classes.
6. Verify that your implementation classes are located in the `${USER_INSTALL_ROOT}/classes` directory so that the WebSphere Application Server class loader can load the classes.

Example

Related tasks

Chapter 5, “Authenticating users,” on page 111

The process of authenticating users involves a user registry and an authentication mechanism. Optionally, you can define trust between WebSphere Application Server and a proxy server, configure single sign-on capability, and specify how to propagate security attributes between application servers.

Using the default single sign-on token with default or custom token factory

Do not use the default single sign-on token in service provider code. This default token is used by the WebSphere Application Server run-time code only.

Before you begin

Size limitations exist for this token when it is added as an HTTP cookie. If you need to create an HTTP cookie using this token framework, you can implement a custom single sign-on token. To implement a custom single sign-on token see “Implementing a custom single sign-on token” on page 806 for more information.

- Modify the single sign-on token factory configuration to use a token factory other than the default token factory.

When the default single sign-on token is generated, the application server utilizes the `TokenFactory` class that is specified using the `com.ibm.wsspi.security.token.singleSignonTokenFactory` property. Use the administrative console to modify the property.

The `com.ibm.ws.security.ltpa.LTPAToken2Factory` token factory is the default that is specified for this property. This token factory creates a single sign-on (SSO) token called `LtpaToken2`, which WebSphere Application Server uses for propagation. This token factory uses the AES/CBC/PKCS5Padding cipher.

If you change this token factory, you lose the interoperability with any servers running a version of WebSphere Application Server prior to version 5.1.1 that use the default token factory. Only servers running WebSphere Application Server Version 5.1.1 or later with propagation enabled are aware of the `LtpaToken2` cookie. If all of your application servers use WebSphere Application Server Version 5.1.1 or later and all of your servers use your new token factory this awareness is not a problem.

1. Open the administrative console.
 2. Click **Security > Global security**.
 3. Under Authentication, click **Custom properties**.
- Perform your own signing and encryption of the default single sign-on token.

If you need to perform your own signing and encryption of the default single sign-on token, you must implement the following classes:

- `com.ibm.wsspi.security.ltpa.Token`
- `com.ibm.wsspi.security.ltpa.TokenFactory`

Your token factory implementation instantiates (`createToken`) and validates (`validateTokenBytes`) your token implementation. You can use the Lightweight Third-Party Authentication (LTPA) keys passed into the `initialize` method of the token factory or you can use your own keys. If you use your own keys, they must be the same everywhere to validate the tokens that are generated using those keys. See the API reference information for more information on implementing your own custom token factory.

- Associate your own token factory with the default single sign-on token.
 1. Open the administrative console.
 2. Click **Security > Global security**.
 3. Under Authentication, click **Custom properties**.
 4. Locate the `com.ibm.wsspi.security.token.singleSignonTokenFactory` property and verify that the value of this property matches your custom `TokenFactory` implementation.
 5. Verify that your implementation classes are put into the `app_server_root/classes` directory so that the WebSphere Application Server class loader can load the classes.
 6. Verify that your implementation classes are located in the `${USER_INSTALL_ROOT}/classes` directory so that the WebSphere Application Server class loader can load the classes.

Related tasks

“Implementing a custom single sign-on token” on page 806

You can create your own single sign-on token implementation. The single sign-on token implementation is set in the login Subject and added to the HTTP response as an HTTP cookie.

Configuring the authentication cache

The security authentication cache affects the frequency of rehashing and the distribution of the hash algorithms.

About this task

To configure the authentication cache properties, complete the following steps:

1. Click **Servers > Application Servers > server_name**.
2. Under Server infrastructure, click **Java and Process Management > Process definition**.
3. Under Additional properties, click **Java Virtual Machine > Custom Properties**.
4. Click **New** to specify a new custom property.

What to do next

For information on the supported authentication cache properties, see “Authentication cache settings” on page 107.

Configuring Common Secure Interoperability Version 2 (CSIV2) inbound and outbound communication settings

WebSphere Application Server enables you to specify Internet Inter-ORB Protocol (IIOP) authentication for both inbound and outbound authentication requests. For inbound requests, you can specify the type of accepted authentication, such as basic authentication. For outbound requests, you can specify properties such as type of authentication, identity assertion or login configurations that are used for requests to downstream servers.

About this task

Complete the following steps to configure Common Secure Interoperability Version 2 (CSIV2) and Security Authentication Service (SAS).

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

1. Determine how to configure security inbound and outbound at each point in your infrastructure. For example, you might have a Java client communicating with an Enterprise JavaBeans (EJB) application server, which in turn communicates to a downstream EJB application server.

The Java client utilizes the `sas.client.props` file to configure outbound security. Pure clients must configure outbound security only.

The upstream EJB application server configures inbound security to handle the correct type of authentication from the Java client. The upstream EJB application server utilizes the outbound security configuration when going to the downstream EJB application server.

This type of authentication might be different than what you expect from the Java client into the upstream EJB application server. Security might be tighter between the pure client and the first EJB server, depending on your infrastructure. The downstream EJB server utilizes the inbound security configuration to accept requests from the upstream EJB server. These two servers require similar configuration options as well. If the downstream EJB application server communicates to other downstream servers, the outbound security might require a special configuration.

2. Specify the type of authentication.

By default, authentication by a user ID and password is performed.

Both Java client certificate authentication and identity assertion are disabled by default. If you want this type of authentication performed at every tier, use the CSIV2 authentication protocol configuration as is. However, if you have any special requirements where some servers authenticate differently from other servers, consider how to configure CSIV2 to its best advantage.

3. Configure clients and servers.

Configuring a pure Java client is done through the `sas.client.props` file, where properties are modified.

Configuring servers is always done from the administrative console or scripting, either from the security navigation for cell-level configurations or from the server security of the application server for server-level configurations. If you want some servers to authenticate differently from others, modify some of the server-level configurations. When you modify the server-level configurations, you are overriding the cell-level configurations.

What to do next

Use CSIV2 inbound communications settings for configuring the type of authentication information that is contained in an incoming request or transport.

Use CSIV2 outbound communications settings to specify the features that a server supports when acting as a client to another downstream server.

Configuring Common Secure Interoperability Version 2 inbound communications

Inbound communications refers to the configuration that determines the type of accepted authentication for inbound requests. This authentication is advertised in the interoperable object reference (IOR) that the client retrieves from the name server.

1. Start the administrative console.
2. Click **Security > Global security**.
3. Under RMI/IIOP security, click **CSIV2 inbound communications**.
4. Consider the following layers of security:
 - Identity assertion (attribute layer).

When selected, this server accepts identity tokens from upstream servers. If the server receives an identity token, the identity is taken from an originating client. For example, the identity is in the same form that the originating client presented to the first server. An upstream server sends the identity of the originating client. The format of the identity can be either a principal name, a distinguished name, or a certificate chain. In some cases, the identity is anonymous. It is important to trust the upstream server that sends the identity token because the identity authenticates on this server. Trust of the upstream server is established either using Secure Sockets Layer (SSL) client certificate

authentication or basic authentication. You must select one of the two layers of authentication in both inbound and outbound authentication when you choose identity assertion.

The server ID is sent in the client authentication token with the identity token. The server ID is checked against the trusted server ID list. If the server ID is on the trusted server list, the server ID is authenticated. If the server ID is valid, the identity token is put into a credential and used for authorization of the request.

For more information, refer to Identity assertion.

- User ID and password (message layer).

This type of authentication is the most typical. The user ID and password or authenticated token is sent from a pure client or from an upstream server. However, the upstream server cannot be a z/OS server because z/OS does not support a user ID or password from a server acting as a client. When a user ID and password are received at the server, they are authenticated with the user registry.

Usually, a token is sent from an upstream server and a user ID and password are sent from a client, including a servlet. When a token is received at the server level, the token is validated to determine whether tampering has occurred or whether it is expired.

For more information, refer to User ID and password.

- Secure Sockets Layer client certificate authentication (transport layer).

The SSL client certificate is used to authenticate instead of using user ID and Password. If a server delegates an identity to a downstream server, the identity comes from either the message layer (a client authentication token) or the attribute layer (an identity token), and not from the transport layer through the client certificate authentication.

A client has an SSL client certificate that is stored in the keystore file of the client configuration. When SSL client authentication is enabled on this server, the server requests that the client send the SSL client certificate when the connection is established. The certificate chain is available on the socket whenever a request is sent to the server. The server request interceptor gets the certificate chain from the socket and maps this certificate chain to a user in the user registry. This type of authentication is optimal for communicating directly from a client to a server. However, when you have to go downstream, the identity typically flows over the message layer or through identity assertion.

5. Consider the following points when deciding what type of authentication to accept:

- A server can receive multiple layers simultaneously, so an order of precedence rule decides which identity to use. The identity assertion layer has the highest priority, the message layer follows, and the transport layer has the lowest priority. The SSL client certificate authentication is used when it is the only layer provided. If the message layer and the transport layer are provided, the message layer is used to establish the identity for authorization. The identity assertion layer is used to establish precedence when provided.
- Does this server usually receive requests from a client, from a server, or both? If the server always receives requests from a client, identity assertion is not needed. You can choose either the message layer, the transport layer, or both. You also can decide when authentication is required or just supported. To select a layer as required, the sending client must supply this layer, or the request is rejected. However, if the layer is only supported, the layer might not be supplied.
- What kind of client identity is supplied? If the client identity is client certificates authentication and you want the certificate chain to flow downstream so that it maps to the downstream server user registries, identity assertion is the appropriate choice. Identity assertion preserves the format of the originating client. If the originating client authenticated with a user ID and password, a principal identity is sent. If authentication is done with a certificate, the certificate chain is sent.

In some cases, if the client authenticated with a token and a Lightweight Directory Access Protocol (LDAP) server is the user registry, then a distinguished name (DN) is sent.

6. Configure a trusted server list. When identity assertion is selected for inbound requests, insert a pipe-separated (|) list of server administrator IDs to which this server can support identity token submission. For backwards compatibility, you can still use a comma-delimited list. However, if the server ID is a distinguished name (DN), then you must use a pipe-delimited (|) list because a comma delimiter does not work. If you choose to support any server sending an identity token, you can enter

an asterisk (*) in this field. This action is called *presumed trust*. In this case, use SSL client certificate authentication between servers to establish the trust.

7. Configure session management. You can choose either *stateful* or *stateless* security. Performance is optimum when choosing stateful sessions. The first method request between a client and server is authenticated. All subsequent requests (or until the credential token expires) reuse the session information, including the credential. A client sends a context ID for subsequent requests. The context ID is scoped to the connection for uniqueness.

Results

When you finish configuring this panel, you have configured most of the information that a client gathers when determining what to send to this server. A client or server outbound configuration with this server inbound configuration, determines the security that is applied. When you know what clients send, the configuration is simple. However, if you have a diverse set of clients with differing security requirements, your server considers various layers of authentication.

For a J2EE application server, the authentication choice is usually either identity assertion or message layer because you want the identity of the originating client delegated downstream. You cannot easily delegate a client certificate using an SSL connection. It is acceptable to enable the transport layer because additional server security, as the additional client certificate portion of the SSL handshake, adds some overhead to the overall SSL connection establishment.

What to do next

After you determine which type of authentication data this server might receive, you can determine what to select for outbound security. For more information, see *Configuring Common Secure Interoperability Version 2* outbound authentication.

Common Secure Interoperability Version 2 inbound communications settings

Use this page to specify the features that a server supports for a client accessing its resources.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. From Authentication, click **RMI/IOP security > CSiv2 inbound communications**.

Use common secure interoperability (CSI) inbound communications settings for configuring the type of authentication information that is contained in an incoming request or transport.

Authentication features include three layers of authentication that you can use simultaneously:

- **CSiv2 attribute layer.** The attribute layer might contain an identity token, which is an identity from an upstream server that already is authenticated. The identity layer has the highest priority, followed by the message layer, and then the transport layer. If a client sends all three, only the identity layer is used. The only way to use the SSL client certificate as the identity is if it is the only information that is presented during the request. The client picks up the interoperable object reference (IOR) from the namespace and reads the values from the tagged component to determine what the server needs for security.
- **CSiv2 transport layer.** The transport layer, which is the lowest layer, might contain a Secure Sockets Layer (SSL) client certificate as the identity.
- **CSiv2 message layer.** The message layer might contain a user ID and password or an authenticated token with an expiration.

Propagate security attributes:

Specifies support for security attribute propagation during login requests. When you select this option, the application server retains additional information about the login request, such as the authentication strength used, and retains the identity and location of the request originator.

If you do not select this option, the application server does not accept any additional login information to propagate to downstream servers.

Default: Enabled

Use identity assertion:

Specifies that identity assertion is a way to assert identities from one server to another during a downstream Enterprise JavaBeans (EJB) invocation.

This server does not authenticate the asserted identity again because it trusts the upstream server. Identity assertion takes precedence over all other types of authentication.

Identity assertion is performed in the attribute layer and is only applicable on servers. The principal determined at the server is based on precedence rules. If identity assertion is used, the identity is always derived from the attribute layer. If basic authentication is used without identity assertion, the identity is always derived from the message layer. Finally, if SSL client certificate authentication is performed without either basic authentication, or identity assertion, then the identity is derived from the transport layer.

The identity asserted is the invocation credential that is determined by the RunAs mode for the enterprise bean. If the RunAs mode is Client, the identity is the client identity. If the RunAs mode is System, the identity is the server identity. If the RunAs mode is Specified, the identity is the one specified. The receiving server receives the identity in an identity token and also receives the sending server identity in a client authentication token. The receiving server validates the sending server identity as a trusted identity through the Trusted Server IDs entry box. Enter a list of pipe-separated (|) principal names, for example, serverid1|serverid2|serverid3.

All identity token types map to the user ID field of the active user registry. For an ITTPPrincipal identity token, this token maps one-to-one with the user ID fields. For an ITTDistinguishedName identity token, the value from the first equal sign is mapped to the user ID field. For an ITTCertChain identity token, the value from the first equal sign of the distinguished name is mapped to the user ID field.

When authenticating to an LDAP user registry, the LDAP filters determine how an identity of type ITTCertChain and ITTDistinguishedName get mapped to the registry. If the token type is ITTPPrincipal, then the principal gets mapped to the UID field in the LDAP registry.

Default: Disabled

Trusted identities:

Specifies the trusted identity that is sent from the sending server to the receiving server.

Use this list to decide whether a server is trusted. Even if the server is on the list, the sending server must still authenticate with the receiving server to accept the identity token of the sending server.

Data type: String

Client certificate authentication:

Specifies that authentication occurs when the initial connection is made between the client and the server during a method request.

In the transport layer, Secure Sockets Layer (SSL) client certificate authentication occurs. In the message layer, basic authentication (user ID and password) is used. Client certificate authentication typically performs better than message layer authentication, but requires some additional setup. These additional steps involve verifying that the server trusts the signer certificate of each client to which it is connected. If the client uses a certificate authority (CA) to create its personal certificate, you only need the CA root certificate in the server signer section of the SSL trust file.

When the certificate is authenticated to a Lightweight Directory Access Protocol (LDAP) user registry, the distinguished name (DN) is mapped based on the filter that is specified when configuring LDAP. When the certificate is authenticated to a local OS user registry, the first attribute of the distinguished name (DN) in the certificate, which is typically the common name, is mapped to the user ID in the registry.

The identity from client certificates is used only if no other layer of authentication is presented to the server.

Never Specifies that clients cannot attempt Secure Sockets Layer (SSL) client certificate authentication with this server.

Supported

Specifies that clients connecting to this server can authenticate using SSL client certificates. However, the server can invoke a method without this type of authentication. For example, anonymous or basic authentication can be used instead.

Required

Specifies that clients connecting to this server must authenticate using SSL client certificates before invoking the method.

Transport:

Specifies whether client processes connect to the server using one of its connected transports.

You can choose either Secure Sockets Layer (SSL), TCP/IP or both as the inbound transport that a server supports. If you specify TCP/IP, the server only supports TCP/IP and cannot accept SSL connections. If you specify SSL-supported, this server can support either TCP/IP or SSL connections. If you specify SSL-required, then any server communicating with this one must use SSL.

Note: This option is not available on the z/OS platform unless both Version 6.0.x and earlier nodes exist in the cell.

TCP/IP

If you select **TCP/IP**, then the server opens a TCP/IP listener port only and all inbound requests do not have SSL protection.

SSL-required

If you select **SSL-required**, then the server opens an SSL listener port only and all inbound requests are received using SSL.

SSL-supported

If you select **SSL-supported**, then the server opens both a TCP/IP and an SSL listener port and most inbound requests are received using SSL.

Provide a fixed port number for the following ports. A zero port number indicates that a dynamic assignment is made at run time.

CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS
CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS
SAS_SSL_SERVERAUTH_LISTENER_ADDRESS

Default: SSL-Supported
Range: TCP/IP, SSL Required, SSL-Supported

SSL settings:

Specifies a list of predefined SSL settings to choose from for inbound connection.

Data type: String
Default: DefaultSSLSettings
DefaultIOPSSL
Range: Any SSL settings configured in the SSL Configuration Repertoire

Message layer authentication:

The following options are available for message layer authentication:

Never Specifies that this server cannot accept user ID and password authentication.

Supported

Specifies that a client communicating with this server can specify a user ID and password. However, a method might be invoked without this type of authentication. For example, an anonymous or client certificate might be used instead.

Required

Specifies that clients communicating with this server must specify a user ID and password for any method request.

Allow client to server authentication with::

Specifies client-to-server authentication using Kerberos, LTPA or Basic authentication.

The following options are available for client to server authentication:

Kerberos (KRB5)

Select to specify Kerberos as the authentication mechanism. You must first configure the Kerberos authentication mechanism. Read about “Configuring Kerberos as the authentication mechanism using the administrative console” on page 266 for more information.

LTPA Select to specify the LTPA token authentication

Basic authentication

Basic authentication is Generic Security Services Username Password (GSSUP). This type of authentication typically involves sending a user ID and a password from the client to the server for authentication.

If you select **Basic Authentication** and **LTPA**, and the active authentication mechanism is LTPA, a user name, password, and LTPA tokens are accepted.

If you select **Basic Authentication** and **KRB5** and the active authentication mechanism is KRB5, a user name, password, Kerberos token and LTPA tokens are accepted.

If you do not select **Basic Authentication**, a user name and password are not accepted by the server.

Login configuration:

Specifies the type of system login configuration to use for inbound authentication.

You can add custom login modules by clicking **Security > Global security**. From Authentication, click **Java Authentication and Authorization Service > System logins**.

Stateful sessions:

Select this option to enable stateful sessions, which are used mostly for performance improvements.

The first contact between a client and server must fully authenticate. However, all subsequent contacts with valid sessions reuse the security information. The client passes a context ID to the server, and the ID is used to look up the session. The context ID is scoped to the connection, which guarantees uniqueness. Whenever the security session is not valid and the authentication retry is enabled, which is the default, the client-side security interceptor invalidates the client-side session and submits the request again without user awareness. This situation might occur if the session does not exist on the server; for example, the server failed and resumed operation. When this value is disabled, each method invocation must authenticate again.

Default: Enabled

Configuring Common Secure Interoperability Version 2 outbound communications

The following choices are available when configuring the Common Secure Interoperability Version 2 (CSlv2) outbound communications panel.

Before you begin

Outbound communications refers to the configuration that determines the type of authentication that is performed for outbound requests to downstream servers. Several *layers* or *methods* of authentication can occur. The downstream server inbound authentication configuration must support at least one choice made in this server outbound authentication configuration. If nothing is supported, the request might go outbound as unauthenticated. This situation does not create a security problem because the authorization runtime is responsible for preventing access to protected resources. However, if you choose to prevent an unauthenticated credential from going outbound, you might want to designate one of the authentication layers as required, rather than supported. If a downstream server does not support authentication, then when authentication is required, the method request fails to go outbound.

About this task

The following choices are available in the Common Secure Interoperability Version 2 (CSlv2) outbound communications panel. Remember that you are not required to complete these steps in the displayed order. Rather, these steps are provided to help you understand your choices for configuring outbound communications.

- Select **Identity Assertion** (attribute layer). When selected, this server sends an identity token to a downstream server if the downstream server supports identity assertion. When an originating client authenticates to this server, the authentication information supplied is preserved in the outbound identity token. If the client authenticating to this server uses client certificate authentication, then the identity token format is a certificate chain, containing the exact client certificate chain from the inbound socket. The same scenario is true for other mechanisms of authentication. Read the Identity Assertion topic for more information.
- Select **User ID** and **Password** (message layer). This type of authentication is the most typical. The user ID and password (if BasicAuth credential) or authenticated token (if authenticated credential) are sent

outbound to the downstream server if the downstream server supports message layer authentication in the inbound authentication panel. Refer to the Message Layer Authentication article for more information.

- Select **SSL Client certificate authentication** (transport layer). The main reason to enable outbound Secure Sockets Layer (SSL) client authentication from one server to a downstream server is to create a trusted environment between those servers. For delegating client credentials, use one of the two layers mentioned previously. However, you might want to create SSL personal certificates for all the servers in your domain, and only trust those servers in your SSL truststore file. No other servers or clients can connect to the servers in your domain, except at the tiers where you want them. This process can protect your enterprise bean servers from access by anything other than your servlet servers.

Configuring session management

About this task

You can choose either *stateful* or *stateless* security. Performance is optimum when choosing stateful sessions. The first method request between this server and the downstream server is authenticated. All subsequent requests reuse the session information, including the credential. A *unique session entry* is defined as the combination of a unique client authentication token and an identity token, scoped to the connection.

Example

Typically, the outbound authentication configuration is for an upstream server to communicate with a downstream server. Most likely, the upstream server is a servlet server and the downstream server is an Enterprise JavaBeans (EJB) server. On a servlet server, the client authentication that is performed to access the servlet can be one of many different types of authentication, including client certificate and basic authentication. When receiving basic authentication data, whether through a prompt login or a form-based login, the basic authentication information is typically authenticated to from a credential of the mechanism type that is supported by the server, such as the Lightweight Third Party Authentication (LTPA). When LTPA is the mechanism, a forwardable token exists in the credential. Choose the message layer (BasicAuth) authentication to propagate the client credentials. If the credential is created using a certificate login and you want to preserve sending the certificate downstream, you might decide to go outbound with identity assertion.

What to do next

Save the configuration and restart the server for the changes to take effect.

Common Secure Interoperability Version 2 outbound communications settings

Use this page to specify the features that a server supports when acting as a client to another downstream server.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. From Authentication, click **RMI/IOP security > CSiv2 outbound communications**.

Authentication features include three layers of authentication that you can use simultaneously:

- **CSiv2 attribute layer**. The attribute layer might contain an identity token, which is an identity from an upstream server that already is authenticated. The identity layer has the highest priority, followed by the message layer, and then the transport layer. If a client sends all three, only the identity layer is used. The only way to use the SSL client certificate as the identity is if it is the only information that is presented during the request. The client picks up the interoperable object reference (IOR) from the namespace and reads the values from the tagged component to determine what the server needs for security.

- **CSlv2 transport layer.** The transport layer, which is the lowest layer, might contain a Secure Sockets Layer (SSL) client certificate as the identity.
- **CSlv2 message layer.** The message layer might contain a user ID and password or an authenticated token with an expiration.

Propagate security attributes:

Specifies to support security attribute propagation during login requests. When you select this option, the application server retains additional information about the login request, such as the authentication strength used, and retains the identity and location of the request originator.

If you do not select this option, the application server does not accept any additional login information to propagate to downstream servers.

Default: Enabled

Use server-trusted identity:

Specifies the server identity that the application server uses to establish trust with the target server. The server identity can be sent using one of the following methods:

- A server ID and password when the server password is specified in the registry configuration.
- A server ID in a Lightweight Third Party Authentication (LTPA) token when the internal server ID is used.

For interoperability with application servers other than WebSphere Application Server, use one of the following methods:

- Configure the server ID and password in the registry.
- Select the **Server-trusted identity** option and specify the trusted identity and password so that an interoperable Generic Security Services Username Password (GSSUP) token is sent instead of an LTPA token.

Default: Disabled

Specify an alternative trusted identity:

Specifies an alternative user as the trusted identity that is sent to the target servers instead of sending the server identity.

This option is recommended for identity assertion. The identity is automatically trusted when it is sent within the same cell and does not need to be in the trusted identities list within the same cell. However, this identity must be in the registry of the target servers in an external cell, and the user ID must be on the trusted identities list or the identity is rejected during trust evaluation.

Default: Disabled

Trusted identity:

Specifies the trusted identity that is sent from the sending server to the receiving server.

If you specify an identity in this field, it can be selected on the panel for your configured user account repository. If you do not specify an identity, a Lightweight Third Party Authentication (LTPA) token is sent between the servers.

Specifies a pipe-separated (|) list of trusted server administrator user IDs, which are trusted to perform identity assertion to this server. For example, serverid1|serverid2|serverid3. The application server

supports the comma (,) character as the list delimiter for backwards compatibility. The application server checks the comma character when the pipe character (|) fails to find a valid trusted server ID.

Use this list to decide whether a server is trusted. Even if the server is on the list, the sending server must still authenticate with the receiving server to accept the identity token of the sending server.

Password:

Specifies the password that is associated with the trusted identity.

Data type: Text

Confirm password:

Confirms the password that is associated with the trusted identity.

Data type: Text

Message layer authentication:

The following options are available for message layer authentication:

Never Specifies that this server cannot accept user ID and password authentication.

Supported

Specifies that a client communicating with this server can specify a user ID and password. However, a method might be invoked without this type of authentication. For example, an anonymous or client certificate might be used instead.

Required

Specifies that clients communicating with this server must specify a user ID and password for any method request.

Allow client to server authentication with::

Specifies client-to-server authentication using Kerberos, LTPA or Basic authentication.

The following options are available for client to server authentication:

Kerberos (KRB5)

Select to specify Kerberos as the authentication mechanism. You must first configure the Kerberos authentication mechanism. Read about “Configuring Kerberos as the authentication mechanism using the administrative console” on page 266 for more information.

LTPA Select to configure and enable Lightweight Third-Party Authentication (LTPA) token authentication.

Basic authentication

Basic authentication is Generic Security Services Username Password (GSSUP). This type of authentication typically involves sending a user ID and a password from the client to the server for authentication.

If you select **Basic Authentication** and **LTPA**, and the active authentication mechanism is LTPA, the server goes with a downstream server with a user name, password or LTPA token.

If you select **Basic Authentication** and **KRB5**, and the active authentication mechanism is KRB5, the server goes with a downstream server with a user name, password, Kerberos token or LTPA token.

If you do not select **Basic Authentication**, the server does not go with a downstream server with a user name and password.

Transport:

Specifies whether client processes connect to the server using one of its connected transports.

You can choose to use either Secure Sockets Layer (SSL), TCP/IP or both as the inbound transport that a server supports. If you specify TCP/IP, the server only supports TCP/IP and cannot accept SSL connections. If you specify SSL-supported, this server can support either TCP/IP or SSL connections. If you specify SSL-required, then any server communicating with this one must use SSL.

Note: This option is not available on the z/OS operating system unless both Version 6.0.x and earlier nodes exist in the cell.

TCP/IP

If you select **TCP/IP**, then the server opens a TCP/IP listener port only and all inbound requests do not have SSL protection.

SSL-required

If you select **SSL-required**, then the server opens an SSL listener port only and all inbound requests are received using SSL.

SSL-supported

If you select **SSL-supported**, then the server opens both a TCP/IP and an SSL listener port and most inbound requests are received using SSL.

Provide a fixed port number for the following ports. A zero port number indicates that a dynamic assignment is made at run time.

CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS
CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS
SAS_SSL_SERVERAUTH_LISTENER_ADDRESS

Default: SSL-Supported
Range: TCP/IP, SSL Required, SSL-Supported

SSL settings:

Specifies a list of predefined SSL settings to choose from for inbound connection.

Data type: String
Default: DefaultSSLSettings
DefaultIOPSSL
Range: Any SSL settings configured in the SSL Configuration Repertoire

Client certificate authentication:

Specifies whether a client certificate from the configured keystore is used to authenticate to the server when the SSL connection is made between this server and a downstream server, provided that the downstream server supports client certificate authentication.

Typically, client certificate authentication has a higher performance than message layer authentication, but requires some additional setup. These additional steps include verifying that this server has a personal certificate and that the downstream server has the signer certificate of this server.

If you select client certificate authentication, the following options are available:

Never Specifies that this server does not attempt Secure Sockets Layer (SSL) client certificate authentication with downstream servers.

Supported

Specifies that this server can use SSL client certificates to authenticate to downstream servers. However, a method can be invoked without this type of authentication. For example, the server can use anonymous or basic authentication instead.

Required

Specifies that this server must use SSL client certificates to authenticate to downstream servers.

Default: Enabled

Login configuration:

Specifies the type of system login configuration to use for inbound authentication.

You can add custom login modules by clicking **Security > Global security**. From Authentication, click **Java Authentication and Authorization Service > System logins**.

Stateful sessions:

Select this option to enable stateful sessions, which are used mostly for performance improvements.

The first contact between a client and server must fully authenticate. However, all subsequent contacts with valid sessions reuse the security information. The client passes a context ID to the server, and the ID is used to look up the session. The context ID is scoped to the connection, which guarantees uniqueness. Whenever the security session is not valid and the authentication retry is enabled, which is the default, the client-side security interceptor invalidates the client-side session and submits the request again without user awareness. This situation might occur if the session does not exist on the server, for example, the server failed and resumed operation. When this value is disabled, every method invocation must authenticate again.

Custom outbound mapping:

Enables the use of custom Remote Method Invocation (RMI) outbound login modules.

The custom login module maps or completes other functions before the predefined RMI outbound call.

To declare a custom outbound mapping, complete the following steps:

1. Click **Security > Global security**.
2. From Authentication, click **Java Authentication and Authorization Service > System logins > New**.

Trusted authentication realms - outbound:

If the RMI/IIOP communication is across different realms, use this link to add outbound trusted realms.

The credential tokens are only sent to the realms that are trusted. In addition, the receiving server should trust this realm using the inbound trusted realms configuration to validate the LTPA token.

Configuring inbound transports

By using this configuration, you can configure a different transport for inbound security versus outbound security.

Before you begin

Inbound transports refer to the types of listener ports and their attributes that are opened to receive requests for this server. Both Common Secure Interoperability Specification, Version 2 (CSlv2) and Secure Authentication Service (SAS) have the ability to configure the transport.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

However, the following differences between the two protocols exist:

- CSlv2 is much more flexible than SAS, which requires Secure Sockets Layer (SSL); CSlv2 does not require SSL.
- SAS does not support SSL client certificate authentication, while CSlv2 does.
- CSlv2 can require SSL connections, while SAS only supports SSL connections.
- SAS always has two listener ports open: TCP/IP and SSL.
- CSlv2 can have as few as one listener port and as many as three listener ports. You can open one port for just TCP/IP or when SSL is required. You can open two ports when SSL is supported, and open three ports when SSL and SSL client certificate authentication is supported.

About this task

Complete the following steps to configure the Inbound transport panels in the administrative console:

1. Click **Security > Global security**.
2. Under RMI/IIOP security, click **CSlv2 inbound communications**.
3. Under Transport, select You can choose to use either Secure Sockets Layer (SSL), TCP/IP or both as the inbound transport that a server supports. If you specify TCP/IP, the server only supports TCP/IP and cannot accept SSL connections. If you specify SSL-supported, this server can support either TCP/IP or SSL connections. If you specify SSL-required, then any server communicating with this one must use SSL.

Note: This option is not available on the z/OS platform unless there are both Version 6.0.x and earlier nodes in the cell.

4. Click **Apply**.
5. Consider fixing the listener ports that you configured.

You complete this action in a different panel, but think about this action now. Most endpoints are managed at a single location, which is why they do not display in the Inbound transport panels. Managing end points at a single location helps you decrease the number of conflicts in your configuration when you assign the endpoints. The location for SSL end points is at each server. The following port names are defined in the End points panel and are used for Object Request Broker (ORB) security:

- CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS - CSlv2 Client Authentication SSL Port
- CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS - CSlv2 SSL Port
- SAS_SSL_SERVERAUTH_LISTENER_ADDRESS - SAS SSL Port
- ORB_LISTENER_PORT - TCP/IP Port

For an application server, click **Servers > Application servers > server_name**. Under Communications, click **Ports**. The Ports panel is displayed for the specified server.

For a node agent, go to **System administration > Node agents > node_name**. Under Additional properties, click **Ports**. The Ports panel for the node agent and deployment manager already are fixed, but you might consider reassigning the ports. For the deployment manager, click **System Administration > Deployment manager**. Under **Additional properties**, click **Ports**.

The Object Request Broker (ORB) on WebSphere Application Server uses a listener port for Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) communications, and is statically specified using configuration dialogs or during migration. If you are working with a firewall, you must

specify a static port for the ORB listener and open that port on the firewall so that communication can pass through the specified port. The endPoint property for setting the ORB listener port is: ORB_LISTENER_ADDRESS.

In the WebSphere Application Server Network Deployment environment, the ORB_LISTENER_ADDRESS end point is specified on the node agent. The location service daemon resides on the node agent and piggybacks onto the ORB listener port, which results in needing the port fixed. Also, you must add the ORB_LISTENER_ADDRESS to the other application servers to set their ORB listener port. Each ORB has a distinct listener port. In WebSphere Application Server Network Deployment, you must specify a different listener port. For example, you might specify the following ports:

- Node agent: ORB_LISTENER_ADDRESS=9000
- Server1: ORB_LISTENER_ADDRESS=9811
- Server2: ORB_LISTENER_ADDRESS=9812

Federated servers can run without the node agent running. When ORB_LISTENER_ADDRESS is set to a value of zero (0) or greater, the server does not depend on the location service daemon to redirect connections to the server. When you set ORB_LISTENER_ADDRESS, all object references in the namespace specify the connection to the server, not the location service daemon. When the server is running without the node agent, all applications must be accessed through the name server that runs on the application server. The client must change the Java Naming Directory Interface (JNDI) reference to use the host and port of the application server.

Table 16.

ORB_LISTENER_ADDRESS

value = 0	The server starts on any available port and does not use the location service daemon.
value > 0	The server starts on the port that is specified by the value you enter. The location service daemon is not used.

Note: Work load management might not work without the node agent running.

Complete the following steps using the administrative console to specify the ORB_LISTENER_ADDRESS port or ports.

Complete the following steps for the node agent and the deployment manager.

- Click **Servers > Application Servers > server_name**. Under Communications, click **Ports > New**.
 - Select **ORB_LISTENER_ADDRESS** from the **Port name** field in the Configuration panel.
 - Enter the IP address, the fully qualified Domain Name System (DNS) host name, or the DNS host name by itself in the **Host** field. For example, if the host name is myhost, the fully qualified DNS name can be myhost.myco.com and the IP address can be 155.123.88.201.
 - Enter the port number in the **Port** field. The port number specifies the port for which the service is configured to accept client requests. The port value is used with the host name. Using the previous example, the port number might be 9000.
- Click **Security > Global security**. Under RMI/IIOP security, click **CSlv2 inbound communications**. Select the SSL settings that are used for inbound requests from CSlv2 clients, and then click **Apply**. Remember that the CSlv2 protocol is used to inter-operate with previous releases. When configuring the keystore and truststore files in the SSL configuration, these files need the right information for inter-operating with previous releases of WebSphere Application Server.

Results

The inbound transport configuration is complete. With this configuration, you can configure a different transport for inbound security versus outbound security. For example, if the application server is the first server that is used by users, the security configuration might be more secure. When requests go to back-end enterprise bean servers, you might lessen the security for performance reasons when you go outbound. With this flexibility you can design the right transport infrastructure to meet your needs.

What to do next

When you finish configuring security, perform the following steps to save, synchronize, and restart the servers:

1. Click **Save** in the administrative console to save any modifications to the configuration.
2. Synchronize the configuration with all node agents.
3. Stop and restart all servers, when synchronized.

Common Secure Interoperability Version 2 transport inbound settings

Use this page to specify which listener ports to open and which Secure Sockets Layer (SSL) settings to use. These specifications determine which transport a client or upstream server uses to communicate with this server for incoming requests.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **RMI/IOP security > CSiv2 inbound transport**.

Transport:

Specifies whether client processes connect to the server using one of its connected transports.

You can choose to use either Secure Sockets Layer (SSL), TCP/IP or both as the inbound transport that a server supports. If you specify TCP/IP, the server only supports TCP/IP and cannot accept SSL connections. If you specify SSL-supported, this server can support either TCP/IP or SSL connections. If you specify SSL-required, then any server communicating with this one must use SSL.

If you specify SSL-supported or SSL-required, decide which set of SSL configuration settings you want to use for the inbound configuration. This decision determines which key file and trust file are used for inbound connections to this server.

TCP/IP

If you select **TCP/IP**, then the server opens a TCP/IP listener port only and all inbound requests do not have SSL protection.

SSL-required

If you select **SSL-required**, then the server opens an SSL listener port only and all inbound requests are received using SSL.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

SSL-supported

If you select **SSL-supported**, then the server opens both a TCP/IP and an SSL listener port and most inbound requests are received using SSL.

Provide a fixed port number for the following ports. A zero port number indicates that a dynamic assignment is made at runtime.

CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS
CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS
SAS_SSL_SERVERAUTH_LISTENER_ADDRESS

Default:

SSL-Supported

Range:

TCP/IP, SSL Required, SSL-Supported

SSL settings:

Specifies a list of predefined SSL settings to choose from for inbound connections.

These settings are configured at the SSL Repertoire panel. To access the SSL Repertoire panel, complete the following steps:

1. Clicking **Security > SSL certificate and key management**.
2. Under configuration settings, click **Manage endpoint security configurations and trust zones**.
3. Expand Inbound and click *inbound_configuration*.
4. Under Related items, click **SSL configurations**.

Data type:	String
Default:	DefaultSSLSettings
DefaultIOPSSL	
Range:	Any SSL settings configured in the SSL Configuration Repertoire

Centrally managed:

Specifies that the selection of an SSL configuration is based upon the outbound topology view for the Java Naming and Directory Interface (JNDI) platform.

Centrally managed configurations support one location to maintain SSL configurations rather than spreading them across the configuration documents.

Default:	Enabled
-----------------	---------

Use specific SSL alias:

Specifies the SSL configuration alias to use for LDAP outbound SSL communications.

This option overrides the centrally managed configuration for the JNDI platform.

z/OS SSL settings:

Specifies a list of predefined Secure Sockets Layer (SSL) settings for inbound connections. Configure these settings on the SSL panel by clicking Secure communications on the administrative console.

Secure Authentication Service inbound transport settings

Use this page to specify transport settings for connections that are accepted by this server using the Secure Authentication Service (SAS) authentication protocol. The SAS protocol is used to communicate securely to enterprise beans with previous releases of the application server.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand RMI/IOP security and click **SAS inbound transport**.

Note: The panel associated with this article displays only when you have a Version 6.0.x server in your environment. SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

SSL Settings:

Specifies a list of predefined SSL settings to choose from for inbound connections.

These settings are configured on the Secure Sockets Layer (SSL) configuration panel. To access the SSL configuration panel, complete the following steps:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations and trust zones**.
2. Expand Inbound > *configuration_name*.
3. Under Related Items, click **SSL configurations**.

Data type: String
Default: DefaultSSLSettings

Configuring outbound transports

By using this configuration, you can configure a different transport for inbound security versus outbound security.

Before you begin

Outbound transports refers to the transport that is used to connect to a downstream server. When you configure the outbound transport, consider the transports that the downstream servers support. If you are considering Secure Sockets Layer (SSL), also consider including the signers of the downstream servers in this server truststore file for the handshake to succeed.

When you select an SSL configuration, that configuration points to keystore and truststore files that contain the necessary signers.

If you configured client certificate authentication for this server by completing the following steps, then the downstream servers contain the signer certificate belonging to the server personal certificate:

1. Click **Security > Global security**.
2. Under RMI/IIOP security, click **CSlv2 outbound communications**.

About this task

Complete the following steps to configure the outbound transport panels.

1. Select the type of transport and the SSL settings by clicking **Security > Global security**. Under RMI/IIOP security, click **CSlv2 outbound communications**. By selecting the type of transport, you choose the transport to use when connecting to downstream servers. The downstream servers support the transport that you choose. If you choose **SSL-Supported**, the transport that is used is negotiated during the connection. If both the client and server support SSL, always select the **SSL-Supported** option unless the request is considered a special request that does not require SSL, such as if an object request broker (ORB) is a request.
2. Select the **SSL required** option if you want to use Secure Sockets Layer communications with the outbound transport.

If you select the **SSL required** option, you can select either the **Centrally managed** or **Use specific SSL alias** option.

Centrally managed

Enables you to specify an SSL configuration for particular scope such as the cell, node, server, or cluster in one location. To use the **Centrally managed** option, you must specify the SSL configuration for the particular set of endpoints. The Manage endpoint security configurations and trust zones panel displays all of the inbound and outbound endpoints that use the SSL protocol. If you expand the Inbound or Outbound section of the panel and click the name of a node, you can specify an SSL configuration that is used for every endpoint on that node. For an outbound transport, you can override the inherited SSL configuration by specifying an SSL configuration for a particular endpoint. To specify an SSL configuration for an outbound

transport, click **Security > SSL certificate and key management > Manage endpoint security configurations and trust zones** and expand **Outbound**.

Use specific SSL alias

Select the **Use specific SSL alias** option if you intend to select one of the SSL configurations in the menu below the option.

This configuration is used only when SSL is enabled for LDAP. The default is *DefaultSSLSettings*. To modify or create a new SSL configuration, complete the steps described in “Creating a Secure Sockets Layer configuration” on page 632.

3. Click **Apply**.

Results

The outbound transport configuration is complete. With this configuration, you can configure a different transport for inbound security versus outbound security. For example, if the application server is the first server used by end users, the security configuration might be more secure. When requests go to back-end enterprise beans servers, you might consider less security for performance reasons when you go outbound. With this flexibility you can design a transport infrastructure that meets your needs.

What to do next

When you finish configuring security, perform the following steps to save, synchronize, and restart the servers.

- Click **Save** in the administrative console to save any modifications to the configuration.
- Synchronize the configuration with all node agents.
- Stop and restart all servers, after synchronization.

Common Secure Interoperability Version 2 outbound transport settings

Use this page to specify which transports and Secure Sockets Layer (SSL) settings this server uses when communicating with downstream servers for outbound requests.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **RMI/IIOP security > CSiv2 outbound transport**.

You also can view this administrative console by completing the following steps:

1. Click **Servers > Application servers > server_name**.
2. Under Security, click **Server security**.
3. Under Additional properties, click **CSiv2 outbound transport**.

Transport:

Specifies whether the client processes connect to the server using one of the server-connected transports.

You can choose to use either SSL, TCP/IP, or Both as the outbound transport that a server supports. If you specify TCP/IP, the server supports only TCP/IP and cannot initiate SSL connections with downstream servers. If you specify SSL-supported, this server can initiate either TCP/IP or SSL connections. If you specify SSL-required, this server must use SSL to initiate connections to downstream servers. When you do specify SSL, decide which set of SSL configuration settings you want to use for the outbound configuration.

This decision determines which keyfile and trustfile to use for outbound connections to downstream servers.

Consider the following options:

TCP/IP

If you select this option, the server opens TCP/IP connections with downstream servers only.

SSL-required

If you select this option, the server opens SSL connections with downstream servers.

SSL-supported

If you select this option, the server opens SSL connections with any downstream server that supports them and opens TCP/IP connections with any downstream servers that do not support SSL.

Default: SSL-supported
Range: TCP/IP, SSL-required, SSL-supported

SSL settings:

Specifies a list of predefined SSL settings for outbound connections. These settings are configured at the SSL Configuration Repertoires panel.

To access the panel, complete the following steps:

1. Click **Security > SSL certificate and key management**.
2. Under Configuration settings, click **Manage endpoint security configurations and trust zones**.
3. Expand Outbound > *outbound_configuration_name*.
4. Under Related items, click **SSL configurations**.

Data type: String
Range: Any SSL settings that are configured in the SSL Configuration Repertoires panel

Note: This field is available only if a Version 6.0.x server exists in your environment.

SSL enabled:

Specifies whether secure socket communication is enabled to the server.

Centrally managed:

Specifies that the selection of an SSL configuration is based upon the outbound topology view for the Java Naming and Directory Interface (JNDI) platform.

Centrally managed configurations support one location to maintain SSL configurations rather than spreading them across the configuration documents.

Default: Enabled

Use specific SSL alias:

Specifies the SSL configuration alias that you want to use for outbound SSL communications.

This option overrides the centrally managed configuration for the JNDI (LDAP) protocol.

Secure Authentication Service outbound transport settings

Use this page to specify transport settings for connections that are accepted by this server using the Secure Authentication Service (SAS) authentication protocol.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand RMI/IOP security and click **SAS outbound transport**.

Note: The panel associated with this article displays only when you have a Version 6.0.x server in your environment.

SSL settings:

Specifies a list of predefined Secure Sockets Layer (SSL) settings to choose from for outbound connections.

These settings are configured on the SSL configuration panel. To access the SSL configuration panel, complete the following steps:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations and trust zones**.
2. Expand Outbound > *configuration_name*.
3. Under Related Items, click **SSL configurations**.

Data type: String
Default: DefaultSSLSettings

Configuring inbound messages

You can use the administrative console to configure inbound messages for CSiv2.

1. In the administrative console, click **Security > Global security**.
2. Under Authentication, expand **RMI/HOP security**.
3. Click **CSiv2 inbound communication**.
4. Optional: Click **Propagate security attributes** or **Use identity assertion**. The **Propagate security attributes** option enables support for security attribute propagation during login requests. When you select this option, the application server retains additional information about the login request, such as the authentication strength used, and retains the identity and location of the request originator.

The **Use identity assertion** option specifies that identity assertion is a way to assert identities from one server to another during a downstream Enterprise JavaBeans (EJB) invocation.

5. Under CSiv2 Message layer authentication, select **Supported**, **Never** or **Required**.

Never Specifies that this server cannot accept an authentication mechanism that you select under **Allow client to server authentication with:**.

Supported

Specifies that clients communicating with this server can specify an authentication mechanism that you select under **Allow client to server authentication with:**. However, a method might be invoked without this type of authentication. For example, an anonymous or client certificate might be used instead.

Required

Specifies that clients communicating with this server must specify an authentication mechanism that you select under **Allow client to server authentication with:**.

6. Under **Allow client to server authentication with:**, select **Kerberos**, **LTPA** and or **Basic authentication**. You can optionally select:

Kerberos

Select to enable authentication using the Kerberos token.

LTPA Select to enable authentication using the Lightweight Third-Party Authentication (LTPA) token.

Basic authentication

This type of authentication typically involves sending a user ID and a password from the client to the server for authentication. This is also known as Generic Security Services Username Password (GSSUP).

This authentication also involves delegating a credential token from an already authenticated credential, provided the credential type is forwardable; for example, LTPA.

If you select **supported** under **CSlv2 Message layer authentication**, and check **KRB5** and **LTPA** under **Allow client to server authentication with:**, then the server does not accept the user name and password.

7. Click **OK**.

Results

You have now configured messages for CSlv2 inbound.

Configuring outbound messages

You can use the administrative console to configure outbound messages for CSlv2.

1. In the administrative console, click **Security > Global security**.
2. Under Authentication, expand **RMI/HOP security**.
3. Click **CSlv2 outbound communication**.
4. Optional: Click **Propagate security attributes** or **Use identity assertion**. The **Propagate security attributes** option enables support for security attribute propagation during login requests. When you select this option, the application server retains additional information about the login request, such as the authentication strength used, and retains the identity and location of the request originator.

The **Use identity assertion** option specifies that identity assertion is a way to assert identities from one server to another during a downstream Enterprise JavaBeans (EJB) invocation.

The **Use server trusted identity** option specifies the server identity that the application server uses to establish trust with the target server.

The **Specify an alternative trusted identity** option enables you to specify an alternative user as the trusted identity that is sent to the target servers instead of sending the server identity. If you select this option you must provide the name of the trusted identity and the password that is associated with the trusted identity.

5. Under CSlv2 Message layer authentication, select **Supported**, **Never** or **Required**.

Never Specifies that this server cannot accept an authentication mechanism that you select under **Allow client to server authentication with:**.

Supported

Specifies that clients communicating with this server can specify an authentication mechanism that you select under **Allow client to server authentication with:**. However, a method might be invoked without this type of authentication. For example, an anonymous or client certificate might be used instead.

Required

Specifies that clients communicating with this server must specify an authentication mechanism that you select under **Allow client to server authentication with:**.

6. Under **Allow client to server authentication with:**, select **Kerberos**, **LTPA** and or **Basic authentication**. You can optionally select:

Kerberos

Select to enable authentication using the Kerberos token.

LTPA Select to enable authentication using the Lightweight Third-Party Authentication (LTPA) token.

Basic authentication

This type of authentication typically involves sending a user ID and a password from the client to the server for authentication. This is also known as Generic Security Services Username Password (GSSUP).

This authentication also involves delegating a credential token from an already authenticated credential, provided the credential type is forwardable; for example, LTPA.

If you select **supported** under **CSlv2 Message layer authentication**, and check **KRB5** and **LTPA** under **Allow client to server authentication with:**, then the server does not accept the user name and password.

7. Optional: Select **Custom outbound mapping**. This option enables the use of custom Remote Method Invocation (RMI) outbound login modules.

Results

You have now configured messages for CSlv2 outbound.

Common Secure Interoperability Version 2 and Security Authentication Service (SAS) client configuration

A secure Java client requires configuration properties to determine how to perform security with a server.

These configuration properties are typically put into a properties file somewhere on the client system and referenced by specifying the following system property on the command line of the Java client. For example, this property accepts any valid Web address.

```
-Dcom.ibm.CORBA.ConfigURL=file:profile_root/properties/sas.client.props
```

When this file is processed by the Object Request Broker (ORB), security can be enabled between the Java client and the target server.

If any syntax problems exist with the ConfigURL property and the `sas.client.props` file is not found, the Java client proceeds to connect insecurely. Errors display indicating the failure to read the ConfigURL property. Typically the problem is related to having two slashes after `file`, which is not valid.

Use the following properties to configure the Secure Authentication Service (SAS) and CSlv2 authentication protocols:

-
- "" on page 492

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Authentication protocol settings for a client configuration

You can use settings in the `sas.client.props` file to configure Security Authentication Service (SAS) and Common Secure Interoperability Version 2 (CSlv2) clients.

Use the following settings in the `app_server_root/properties/sas.client.props` file to configure SAS and CSlv2 clients.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Note: The `sas.client.props` file for WebSphere Application Server Version 7.0 now contains some new properties that support BasicAuth and Kerberos, such as:

```

com.ibm.IPC.authenticationTarget=BasicAuth
com.ibm.IPC.loginUserId=
com.ibm.IPC.loginPassword=
com.ibm.IPC.loginSource=prompt
com.ibm.IPC.krb5Service=WAS
com.ibm.IPC.krb5CcacheFile=
com.ibm.IPC.krb5ConfigFile=

```

Related tasks

“Configuring Common Secure Interoperability Version 2 (CSIV2) inbound and outbound communication settings” on page 466

WebSphere Application Server enables you to specify Internet Inter-ORB Protocol (IIOP) authentication for both inbound and outbound authentication requests. For inbound requests, you can specify the type of accepted authentication, such as basic authentication. For outbound requests, you can specify properties such as type of authentication, identity assertion or login configurations that are used for requests to downstream servers.

Related reference

“Common Secure Interoperability Version 2 and Security Authentication Service (SAS) client configuration” on page 488

A secure Java client requires configuration properties to determine how to perform security with a server.

com.ibm.CORBA.securityEnabled:

Use to determine if security is enabled for the client process.

Setting	Value
Data Type	Boolean
Default	True
Valid values	True or false

com.ibm.CSI.protocol:

Use to determine which authentication protocols are active.

The client can configure protocols of `ibm`, `csiv2` or both as active. The only possible values for an authentication protocol are `ibm`, `csiv2` and `both`. Do not use `sas` for the value of an authentication protocol. This restriction applies to both client and server configurations. The following list provides information about using each of these protocol options:

ibm Use this authentication protocol option when you are communicating with WebSphere Application Server Version 4.x or earlier servers.

csiv2 Use this authentication protocol option when you are communicating with WebSphere Application Server Version 5 or later servers because the SAS interceptors are not loaded and running for each method request.

both Use this authentication protocol option for interoperability between WebSphere Application Server Version 4.x or earlier servers and WebSphere Application Server Version 5 or later servers. Typically, specifying both provides greater interoperability with other servers.

Setting	Value
Data type	String
Default	Both
Valid values	<code>ibm</code> , <code>csiv2</code> , <code>both</code>

com.ibm.CORBA.authenticationTarget:

Use to determine the type of authentication mechanism for sending security information from the client to the server.

If basic authentication is specified, the user ID and password are sent to the server. Using the Secure Sockets Layer (SSL) transport with this type of authentication is recommended; otherwise, the password is not encrypted. The target server must support the specified authentication target.

If you specify Lightweight Third Party Authentication (LTPA), then LTPA must be the mechanism configured at the server for a method request to proceed securely.

Setting	Value
Data type	String
Default	BasicAuth
Valid values	BasicAuth, LTPA

com.ibm.CORBA.validateBasicAuth:

Use to determine if the user ID and password get validated immediately after the login data is entered when the authenticationTarget property is set to BasicAuth.

In previous releases, BasicAuth logins validated only with the initial method request. During the first request, the user ID and password are sent to the server. This request is the first time that the client can notice an error, if the user ID or password is incorrect. The validateBasicAuth method is specified and the validation of the user ID and password occurs immediately to the security server.

For performance reasons, you might want to disable this property if you do not want to verify the user ID and password immediately. If the client program can wait, it is better to have the initial method request flow to the user ID and password. However, program logic might not be this simple because of error handling considerations.

Setting	Value
Data type	Boolean
Default	True
Valid values	True, False

com.ibm.CORBA.authenticationRetryEnabled:

Use to specify that a failed login attempt is retried. This property determines if a retry occurs for other errors, such as stateful sessions that are not found on a server or validation failures at the server because of an expiring credential.

The minor code in the exception that is returned to a client determines which errors are retried. The number of retry attempts is dependent upon the com.ibm.CORBA.authenticationRetryCount property.

Setting	Value
Data type	Boolean
Default	True
Valid values	True, False

com.ibm.CORBA.authenticationRetryCount:

Use to specify the number of retries that occur until either a successful authentication occurs or the maximum retry value is reached.

When the maximum retry value is reached, the authentication exception is returned to the client.

Setting	Value
Data type	Integer
Default	3
Range	1-10

com.ibm.CORBA.loginSource:

Use to specify how the request interceptor attempts to log in if it does not find an invocation credential already set.

This property is valid only if message layer authentication occurs. If only transport layer authentication occurs, this property is ignored. When specifying properties, the following two additional properties must be defined:

- `com.ibm.CORBA.loginUserId`
- `com.ibm.CORBA.loginPassword`

When performing a programmatic login, it is not necessary to specify none as the login source. The request fails if a credential is set as the invocation credential during a method request.

Note: For the distributed platform, you can choose to NOT edit the properties file, `sas.client.props`, but set the `loginSource` property as follows: `com.ibm.CORBA.loginSource=none`

When you set `com.ibm.CORBA.loginSource=none` for a remote method invocation (RMI) connection, whether using scripting with `wsadmin` or from other clients, you have to perform a programmatic logon because the logged-in user's credentials are not inherited. You must specify user and/or password at the command line.

Setting	Value
Data type	String
Default	Prompt
Valid values	Prompt, key file, stdin, none, properties

com.ibm.CORBA.loginUserId:

Use to specify the user ID when a properties login is configured and message layer authentication occurs.

This property is valid only when `com.ibm.CORBA.loginSource=properties`. Also set the `com.ibm.CORBA.loginPassword` property.

Setting	Value
Data type	String
Range	Any string that is appropriate for a user ID in the configured user registry of the server.

com.ibm.CORBA.loginPassword:

Use to specify the password when a properties login is configured and message layer authentication occurs.

This property is valid only when `com.ibm.CORBA.loginSource=properties`. Also set the `com.ibm.CORBA.loginuserid` property.

Setting	Value
Data type	String
Range	Any string that is appropriate for a password in the configured user registry of the server.

com.ibm.CORBA.keyFileName:

Use to specify the key file that is used to log in.

A key file is a file that contains a list of realm, user ID, and password combinations that a client uses to log into multiple realms. The realm that is used is the one found in the interoperable object reference (IOR) for the current method request. The value of this property is used when the `com.ibm.CORBA.loginSource=key` file is used.

Setting	Value
Data type	String
Default	C:/WebSphere/AppServer/properties/wssserver.key
Range	Any fully qualified path and file name of a WebSphere Application Server key file.

com.ibm.CORBA.loginTimeout:

Use to specify the length of time that the login prompt stays available before it is considered a failed login.

Setting	Value
Data type	Integer
Units	Seconds
Default	300 (5 minute intervals)
Range	0 - 600 (10 minute intervals)

com.ibm.CORBA.securityEnabled:

Use to determine if security is enabled for the client process.

Setting	Value
Data type	Boolean
Default	True
Range	True, False

Security Authentication Service authentication protocol client settings:

In addition to those properties which are valid for both Security Authentication Service (SAS) and Common Secure Interoperability Version 2 (CSIv2), this article documents properties which are valid only for the SAS authentication protocol.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

com.ibm.CORBA.standardPerformQOPModels:

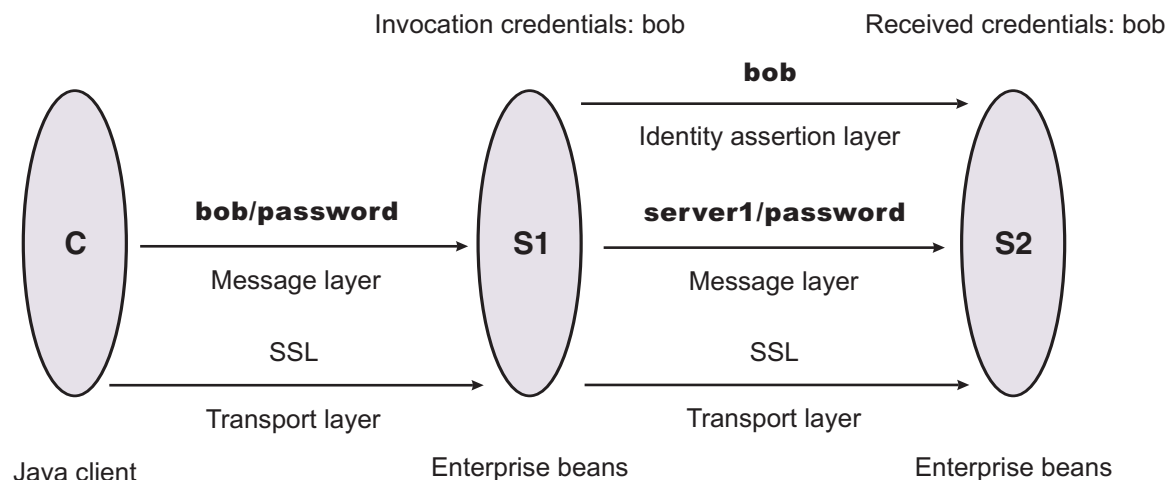
Specifies the strength of the ciphers when making a Secure Sockets Layer (SSL) connection.

Data type:	String
Default:	High
Range:	Low, Medium, High

Example 1: Configuring basic authentication and identity assertion

This example presents a pure Java client, C, that accesses a secure enterprise bean on server, S1, through user bob. The following steps take you through the configuration of C, S1, and S2.

About this task



The enterprise bean code on S1 accesses another enterprise bean on server, S2. This configuration uses identity assertion to propagate the identity of bob to the downstream server, S2. S2 trusts that bob already is authenticated by S1 because it trusts S1. To gain this trust, the identity of S1 also flows to S2 simultaneously and S2 validates the identity by checking the trustedPrincipalList list to verify that it is a valid server principal. S2 also authenticates S1.

1. Configure the client C for message layer authentication with a Secure Sockets Layer (SSL) transport.
 - a. Point the client to the sas.client.props file.
Use the `com.ibm.CORBA.ConfigURL=file:/C:/was/properties/sas.client.props` property. All further configuration involves setting properties within this file.
 - b. Enable SSL.
In this case, SSL is supported but not required:
`com.ibm.CSI.performTransportAssocSSLTLSSupported=true,`
`com.ibm.CSI.performTransportAssocSSLTLSRequired=false`
 - c. Enable client authentication at the message layer.
In this case, client authentication is supported but not required:
`com.ibm.CSI.performClientAuthenticationRequired=false,`
`com.ibm.CSI.performClientAuthenticationSupported=true`
 - d. Use all of the remaining defaults in the sas.client.props file.
2. Configure the server, S1.

In the administrative console, server S1 is configured for incoming requests to support message-layer client authentication and incoming connections to support SSL without client certificate authentication. Server S1 is configured for outgoing requests to support identity assertion.

- a. Configure S1 for incoming connections.
 - 1) Disable identity assertion.
 - 2) Enable user ID and password authentication.
 - 3) Enable SSL.
 - 4) Disable SSL client certificate authentication.
- b. Configure S1 for outgoing connections.
 - 1) Enable identity assertion.
 - 2) Disable user ID and password authentication.
 - 3) Enable SSL.
 - 4) Disable SSL client certificate authentication.
3. Configure the server, S2.

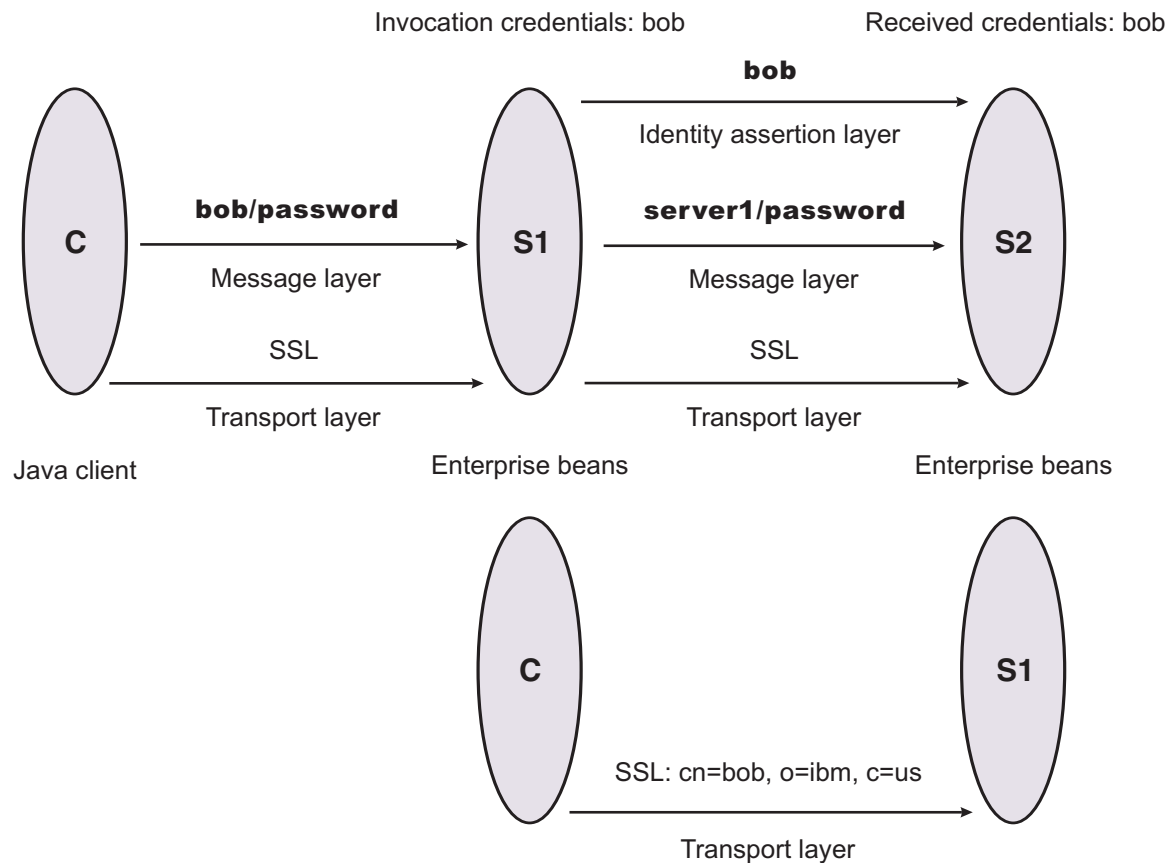
In the administrative console, server S2 is configured for incoming requests to support identity assertion and to accept SSL connections. Complete the following steps to configure incoming connections. Configuration for outgoing requests and connections are not relevant for this example.

- a. Enable identity assertion.
- b. Disable user ID and password authentication.
- c. Enable SSL.
- d. Disable SSL client authentication.

Example 2: Configuring basic authentication, identity assertion, and client certificates

This example is the same as example 1, except for the interaction from client C2 to server S2. Therefore, the configuration of example 1 still is valid, but you have to modify server S2 slightly and add a configuration for client C2. The configuration is not modified for C1 or S1.

About this task



1. Configure client C2 for transport layer authentication (Secure Sockets Layer (SSL) client certificates).

- Point the client to the sas.client.props file.

Use the `com.ibm.CORBA.ConfigURL=file:/C:/was/properties/sas.client.props` property. All further configuration involves setting properties within this file.

- Enable SSL.

In this case, SSL is supported but not required:

```
com.ibm.CSI.performTransportAssocSSLTLSSupported=true,
com.ibm.CSI.performTransportAssocSSLTLSRequired=false
```

- Disable client authentication at the message layer.

```
com.ibm.CSI.performClientAuthenticationRequired=false,
com.ibm.CSI.performClientAuthenticationSupported=false
```

- Enable client authentication at the transport layer where it is supported, but not required.

```
com.ibm.CSI.performTLClientAuthenticationRequired=false,
com.ibm.CSI.performTLClientAuthenticationSupported=true
```

2. Configure the server, S2.

In the administrative console, server S2 is configured for incoming requests to SSL client authentication and identity assertion. Configuration for outgoing requests is not relevant for this example.

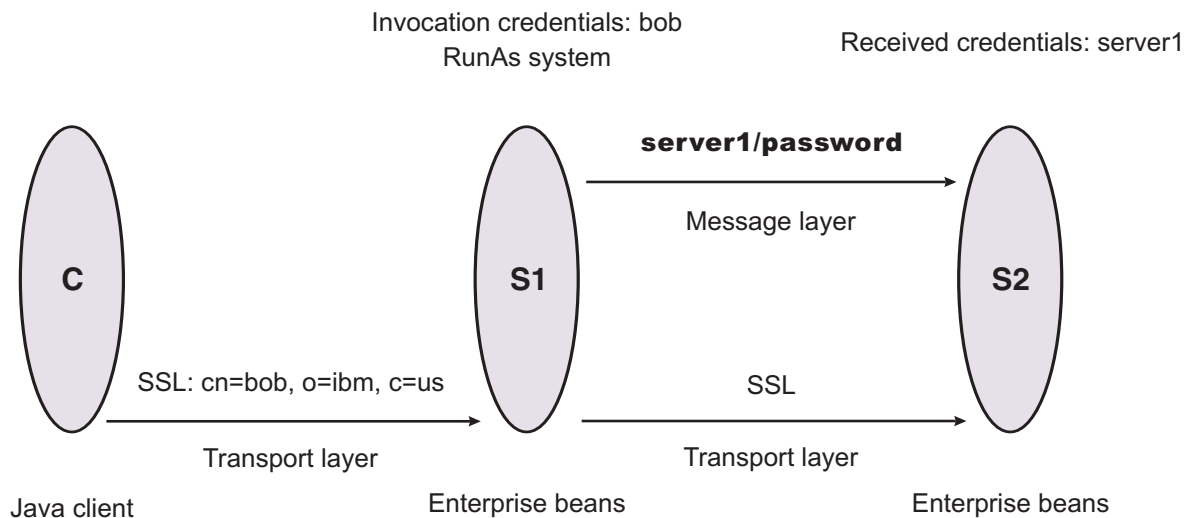
You can mix and match these configuration options. However, a precedence exists as to which authentication features become the identity in the received credential:

- a. Identity assertion
- b. Message-layer client authentication (basic authentication or token)
- c. Transport-layer client authentication (SSL certificates)
- a. Enable identity assertion.
- b. Disable user ID and password authentication.
- c. Enable SSL.
- d. Enable SSL client authentication.

Example 3: Configuring client certificate authentication and RunAs system

This example presents a pure Java client, C, accessing a secure enterprise bean on S1.

About this task



C authenticates to S1 using Secure Sockets Layer (SSL) client certificates. S1 maps the common name of the distinguished name (DN) in the certificate to a user in the local registry. The user in this case is bob. The enterprise bean code on S1 accesses another enterprise bean on S2. Because the RunAs mode is system, the invocation credential is set as server1 for any outbound requests.

1. Configure client C for transport layer authentication (SSL client certificates).
 - a. Point the client to the sas.client.props file.
Use the `com.ibm.CORBA.ConfigURL=file:/C:/was/properties/sas.client.props` property. All further configuration involves setting properties within this file.
 - b. Enable SSL.
In this case, SSL is supported but not required:
`com.ibm.CSI.performTransportAssocSSLTLSSupported=true,`
`com.ibm.CSI.performTransportAssocSSLTLSRequired=false`
 - c. Disable client authentication at the message layer.
`com.ibm.CSI.performClientAuthenticationRequired=false,`
`com.ibm.CSI.performClientAuthenticationSupported=false`
 - d. Enable client authentication at the transport layer. It is supported, but not required.
`com.ibm.CSI.performTLClientAuthenticationRequired=false,`
`com.ibm.CSI.performTLClientAuthenticationSupported=true`

2. Configure the S1 server. In the administrative console, S1 is configured for incoming connections to support SSL with client certificate authentication. The S1 server is configured for outgoing requests to support message layer client authentication.
 - a. Configure S1 for incoming connections.
 - 1) Disable identity assertion.
 - 2) Disable user ID and password authentication.
 - 3) Enable SSL.
 - 4) Enable SSL client certificate authentication.
 - b. Configure S1 for outgoing connections.
 - 1) Disable identity assertion.
 - 2) Disable user ID and password authentication.
 - 3) Enable SSL.
 - 4) Enable SSL client certificate authentication.
3. Configure the S2 server.


In the administrative console, the S2 server is configured for incoming requests to support message layer authentication over SSL. Configuration for outgoing requests is not relevant for this scenario.

 - a. Disable identity assertion.
 - b. Enable user ID and password authentication.
 - c. Enable SSL.
 - d. Disable SSL client authentication.

Example 4: Configuring TCP/IP transport using a virtual private network

This scenario illustrates the ability to choose TCP/IP as the transport when it is appropriate. In some cases, when two servers are on the same virtual private network (VPN), it can be appropriate to select TCP/IP as the transport for performance reasons because the VPN already encrypts the message.

About this task



1. Configure client C for message layer authentication with an Secure Sockets Layer (SSL) transport.
 - a. Point the client to the `sas.client.props` file.
Use the `com.ibm.CORBA.ConfigURL=file:/C:/was/properties/sas.client.props` property. All further configuration involves setting properties within this file.
 - b. Enable SSL.
In this case, SSL is supported but not required.
`com.ibm.CSI.performTransportAssocSSLTLSSupported=true,`
`com.ibm.CSI.performTransportAssocSSLTLSRequired=false`
 - c. Enable client authentication at the message layer. In this case, client authentication is supported but not required. `com.ibm.CSI.performClientAuthenticationRequired=false,`
`com.ibm.CSI.performClientAuthenticationSupported=true`
 - d. Use the remaining defaults in the `sas.client.props` file.
2. Configure the S1 server. In the administrative console, the S1 server is configured for incoming requests to support message-layer client authentication and incoming connections to support SSL without client certificate authentication. The S1 server is configured for outgoing requests to support identity assertion.
It is possible to enable SSL for inbound connections and disable SSL for outbound connections. The same is true in reverse.
 - a. Configure S1 for incoming connections.
 - 1) Disable identity assertion.

- 2) Enable user ID and password authentication.
- 3) Enable SSL.
- 4) Disable SSL client certificate authentication.
- b. Configure S1 for outgoing connections.
 - 1) Disable identity assertion.
 - 2) Enable user ID and password authentication.
 - 3) Disable SSL.
3. Configure the S2 server.

In the administrative console, the S2 server is configured for incoming requests to support identity assertion and to accept SSL connections. Configuration for outgoing requests and connections are not relevant for this scenario.

- a. Disable identity assertion.
- b. Enable user ID and password authentication.
- c. Disable SSL.

Authentication protocol for EJB security

WebSphere Application Server Version 7.0 servers support the CSiv2 authentication protocol only. SAS is only supported between Version 6.0.x and earlier version servers that have been federated in a Version 7.0 cell. The option to select between SAS, CSiv2, or both is only available in the administration console when a Version 6.0.x or earlier release has been federated in a Version 7.0 cell.

SAS is the authentication protocol used by all previous releases of WebSphere Application Server and is maintained for backwards compatibility. The Object Management Group (OMG) has defined the authentication protocol called CSiv2 so that vendors can interoperate securely. CSiv2 is implemented in WebSphere Application Server with more features than SAS and is considered the strategic protocol.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Invoking Enterprise Java Beans (EJB) methods in a secure WebSphere Application Server environment requires an authentication protocol to determine the level of security and the type of authentication that occur between any given client and server for each request. It is the job of the authentication protocol during a method invocation to merge the server authentication requirements that are determined by the object Interoperable Object Reference (IOR) with the client authentication requirements that are determined by the client configuration and come up with an authentication policy specific to that client and server pair.

The authentication policy makes the following decisions, among others, which are all based on the client and server configurations:

- What kind of connection can you make to this server--Secure Sockets Layer (SSL) or TCP/IP?
- If SSL is chosen, how strong is the encryption of the data?
- If SSL is chosen, do you authenticate the client using client certificates?
- Do you authenticate the client with a user ID and password? Does an existing credential exist?
- Do you assert the client identity to downstream servers?
- Given the configuration of the client and server, can a secure request proceed?

You can configure both protocols (SAS and CSiv2) to work simultaneously. If a server supports both protocols, it exports an IOR containing tagged components describing the configuration for SAS and CSiv2. If a client supports both protocols, it reads tagged components for both CSiv2 and SAS. If the client supports both and the server supports both, CSiv2 is used. However, if the server supports SAS (for example, it is a previous WebSphere Application Server release) and the client supports both, the client chooses SAS for this request because the SAS protocol is what both have in common.

Choose a protocol by specifying the `com.ibm.CSI.protocol` property on the client side and configuring through the administrative console on the server side. More details are included in the SAS and CSiv2 properties articles.

Common Secure Interoperability Specification, Version 2

The Common Secure Interoperability Specification, Version 2 (CSiv2) defines the Security Attribute Service (SAS) that enables interoperable authentication, delegation, and privileges. The CSiv2 SAS and SAS protocols are entirely different. The CSiv2 SAS is a subcomponent of CSiv2 that supports SSL and interoperability with the EJB Specification, Version 2.1.

Security Attribute Service

The Common Secure Interoperability Specification, Version 2 Security Attribute Service (CSiv2 SAS) protocol is designed to exchange its protocol elements in the service context of a General Inter-ORB Protocol (GIOP) request and reply messages that are communicated over a connection-based transport. The protocol is intended for use in environments where transport layer security, such as that available through Secure Sockets Layer (SSL) and Transport Layer Security (TLS), is used to provide message protection (that is, integrity and or confidentiality) and server-to-client authentication. The protocol provides client authentication, delegation, and privilege functionality that might be applied to overcome corresponding deficiencies in an underlying transport. The CSiv2 SAS protocol facilitates interoperability by serving as the higher-level protocol under which secure transports can be unified.

Connection and request interceptors

The authentication protocols that are used by WebSphere Application Server are add-on Interoperable Inter-ORB Protocol (IIOP) services. IIOP is a request-and-reply communications protocol that is used to send messages between two Object Request Brokers (ORBs). For each request made by a client ORB to a server ORB, an associated reply is made by the server ORB back to the client ORB. Prior to any request flowing, a connection between the client ORB and the server ORB must be established over the TCP/IP transport (SSL is a secure version of TCP/IP). The client ORB invokes the authentication protocol client connection interceptor, which is used to read the tagged components in the IOR of the object that is located on the server. As mentioned previously, the authentication policy is established here for the request. Given the authentication policy (a coalescing of the server configuration with the client configuration), the strength of the connection is returned to the ORB. The ORB makes the appropriate connection, usually over SSL.

After the connection is established, the client ORB invokes the authentication protocol client request interceptor, which is used to send security information other than what is established by the transport. The security information includes the user ID and password token that are authenticated by the server, an authentication mechanism-specific token that is validated by the server, or an identity assertion token. Identity assertion is a way for one server to trust another server without the need to re-authenticate or re-validate the originating client. However, some work is required for the server to trust the upstream server. This additional security information is sent with the message in a *service context*. A service context has a registered identifier so that the server ORB can identify which protocol is sending the information.

The fact that a service context contains a unique identity is another way for WebSphere Application Server to support both SAS and CSiv2 simultaneously because both protocols have different service context IDs. After the client request interceptor finishes adding the service context to the message, the message is sent to the server ORB.

When the message is received by the server ORB, the ORB invokes the authentication protocol server request interceptor. This interceptor looks for the service context ID known by the protocol. When both SAS and CSiv2 are supported by a server, two different server request interceptors are invoked and both interceptors look for different service context IDs.

However, only one finds a service context for any given request. When the server request interceptor finds a service context, it reads the information in the service context. A method is invoked to the security server to authenticate or validate client identity. The security server either rejects the information or returns a credential. A credential contains additional information about the client that is retrieved from the user registry so that authorization can make the appropriate decision. Authorization is the process of determining if the user can invoke the request based on the roles that are applied to the method and the roles given to the user.

If a service context is not found by the CSIv2 server request interceptor, the interceptor process looks at the transport connection to see if a client certificate chain is sent. This process is done when SSL client authentication is configured between the client and server.

If a client certificate chain is found, the distinguished name (DN) is extracted from the certificate and is used to map to an identity in the user registry. If the user registry is Lightweight Directory Access Protocol (LDAP), the search filters defined in the LDAP registry configuration determine how the certificate maps to an entry in the registry. If the user registry is local OS, the first attribute of the distinguished name (DN) maps to the user ID of the registry. This attribute is typically the common name.

If the certificate does not map, no credential is created and the request is rejected. When valid security information is not presented, the method request is rejected and a `NO_PERMISSION` exception is sent back with the reply. However, when no security information is presented, an unauthenticated credential is created for the request and the authorization engine determines if the method gets invoked. For an unauthenticated credential to invoke an Enterprise JavaBean (EJB) method, either no security roles are defined for the method or a special Everyone role is defined for the method.

When the method invocation is completed in the EJB container, the server request interceptor is invoked again to complete server authentication and a new reply service context is created to inform the client request interceptor of the outcome. This process is typically for making the request *stateful*. When a stateful request is made, only the first request between a client and server requires that security information is sent. All subsequent method requests need to send a unique context ID only so that the server can look up the credential that is stored in a session table. The context ID is unique within the connection between a client and server.

Finally, the method request cycle is completed by the client request interceptor receiving a reply from the server with a reply service context providing information so that the client-side stateful context ID can be confirmed and reused.

Specifying a stateful client is done through the property `com.ibm.CSI.performStateful` (true/false). Specifying a stateful server is done through the administrative console configuration.

Authentication protocol flow

Step 1:

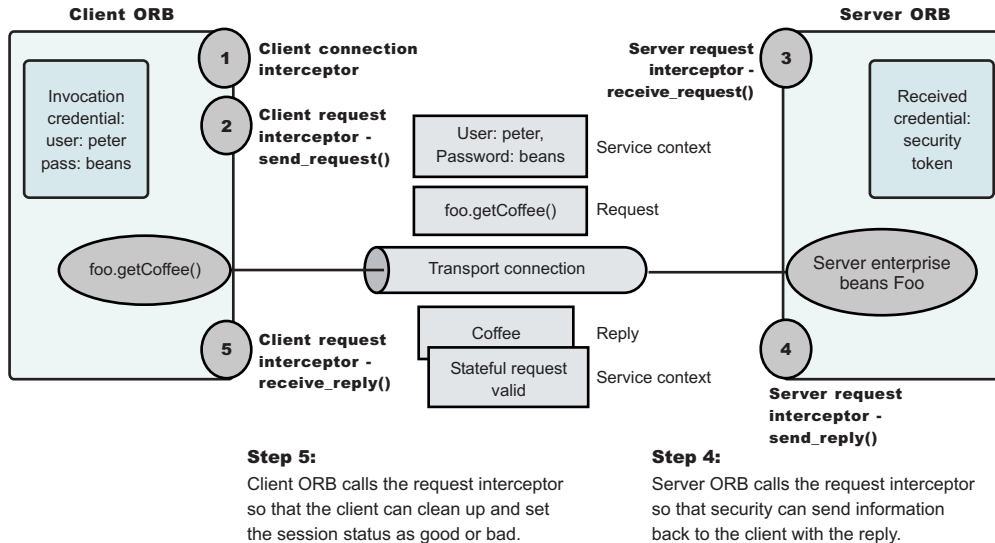
Client ORB calls the connection interceptor to create the connection.

Step 2:

Client ORB calls the request interceptor to get client security information.

Step 3:

Server ORB calls the request interceptor to receive the security information, authenticate, and set the received credential.



. Authentication protocol flow

Authentication policy for each request

The authentication policy of a given request determines the security protection between a client and a server. A client or server authentication protocol configuration can describe required features, supported features, and non-supported features. When a client requires a feature, it can talk only to servers that either require or support that feature. When a server requires a feature, it can talk only to clients that either require or support that feature. When a client supports a feature, it can talk to a server that supports or requires that feature, but can also talk to servers that do not support the feature. When a server supports a feature, it can talk to a client that supports or requires the feature, but can also talk to clients that do not support the feature or chose not to support the feature.

For example, for a client to support client certificate authentication, some setup is required to either generate a self-signed certificate or to get one from a certificate authority (CA). Some clients might not need to complete these actions, therefore, you can configure this feature as not supported. By making this decision, the client cannot communicate with a secure server that requires client certificate authentication. Instead, this client can choose to use the user ID and password as the method of authenticating itself to the server.

Typically, supporting a feature is the most common way of configuring features. It is also the most successful during runtime because it is more forgiving than requiring a feature. Knowing how secure servers are configured in your domain, you can choose the right combination for the client to ensure successful method invocations and still get the most security. If you know that all of your servers support both client certificate and user ID and password authentication for the client, you might want to require one and not support the other. If both the user ID and password and the client certificate are supported on the client and server, both are performed, but user ID and password take precedence at the server. This action is based on the CSiv2 specification requirements.

Authentication protocol support

Use this page to reference information regarding supported authentication protocols.

Authentication protocol support

Beginning with WebSphere Application Server Version 7.0, the WebSphere Application Server Version 7.0 servers only support the Common Secure Interoperability Version 2 (CSIv2) authentication protocol. Secure Authentication Service (SAS) is only supported between Version 6.0.x and previous version servers that have been federated in a Version 7.0 cell. The option to select between SAS, CSIv2, or both will only be made available in the administration console when a Version 6.0.x or previous release has been federated in a Version 7.0 cell.

In future releases, IBM will no longer ship or support the Secure Authentication Service (SAS) IIOP security protocol. It is recommended that you use the Common Secure Interoperability version 2 (CSIv2) protocol.

You can configure both protocols to work simultaneously between Version 6.0.x and previous version servers that have been federated in a Version 7.0 cell. If a server supports both protocols, it exports an interoperable object reference (IOR) that contains tagged components describing the configuration for SAS and CSIv2. If a client supports both protocols, it reads tagged components for both CSIv2 and SAS. If the client and server support both protocols, CSIv2 is used. However, if the server supports SAS (for example, the server is a previous WebSphere Application Server release) and the client supports both protocols, the client chooses SAS for this request.

Choose a protocol using the `com.ibm.CSI.protocol` property on the client side and configure this protocol through the administrative console on the server side.

Common Secure Interoperability Version 2 features

The following Common Secure Interoperability Version 2 (CSIv2) features are available in IBM WebSphere Application Server: message layer authentication, identity assertion, and security attribute propagation.

- Identity Assertion

Supports a downstream server in accepting the client identity that is established on an upstream server, without having to authenticate again. The downstream server trusts the upstream server.

- Message Layer Authentication

Authenticates credential information and sends that information across the network so that a receiving server can interpret it.

- Security attribute propagation

Supports the use of the authorization token to propagate serialized Subject contents and PropagationToken contents with the request. You can propagate these objects using a pure client or a server login that adds custom objects to the Subject. Propagating security attributes prevents downstream logins from having to make user registry calls to look up these attributes.

Propagating security attributes is also useful when the security attributes contain information that is only available at the time of authentication. This information cannot be located using the user registry on downstream servers.

Identity assertion to the downstream server

When a client authenticates to a server, the received credential is set. When the authorization engine checks the credential to determine whether access is permitted, it also sets the *invocation* credential. *Identity assertion* is the invocation credential that is asserted to the downstream server.

When a client authenticates to a server, the received credential is set. When the authorization engine checks the credential to determine whether access is permitted, it also sets the *invocation* credential so that if the Enterprise JavaBeans (EJB) method calls another EJB method that is located on other servers, the invocation credential can be the identity used to invoke the downstream method. Depending on the RunAs mode for the enterprise beans, the invocation credential is set as the originating client identity, the server identity, or a specified different identity. Regardless of the identity that is set, when identity assertion is enabled, it is the invocation credential that is asserted to the downstream server.

The invocation credential identity is sent to the downstream server in an identity token. In addition, the sending server identity, including the password or token, is sent in the client authentication token when basic authentication is enabled. The sending server identity is sent through a Secure Sockets Layer (SSL) client certification authentication when client certificate authentication is enabled. Basic authentication takes precedence over client certificate authentication.

Both identity tokens are needed by the receiving server to accept the asserted identity. The receiving server completes the following actions to accept the asserted identity:

- The server determines whether the sending server identity, sent with a basic authentication token or with an SSL client certificate, is on the trusted principal list of the receiving server. The server determines whether the sending server can send an identity token to the receiving server.
- After it is determined that the sending server is on the trusted list, the server authenticates the sending server to verify its identity.
- The server is authenticated by comparing the user ID and password from the sending server to the receiving server. If the credentials of the sending server are authenticated and on the trusted principal list, then the server proceeds to evaluate the identity token.
- The downstream server checks its defined user registry for the presence of the asserted user ID to gather additional credential information for authorization purposes (for example, group memberships). Thus, the downstream user registry must contain all of the asserted user IDs. Otherwise, identity assertion is not possible. In a stateful server, this action occurs once for the sending server and the receiving server pair where the identity tokens are the same. Subsequent requests are made through a session ID.

Note: When the downstream server does not have a user registry with access to the asserted user IDs in its repository, do not use identity assertion because authorization checks will fail. By disabling identity assertion, the authorization checks on the downstream server are not needed.

Evaluation of the identity token consists of the following four identity formats that exist in an identity token:

- Principal name
- Distinguished name
- Certificate chain
- Anonymous identity

The product servers that receive authentication information typically support all four identity types. The sending server decides which one is chosen, based on how the original client authenticated. The existing type depends on how the client originally authenticates to the sending server. For example, if the client uses Secure Sockets Layer (SSL) client authentication to authenticate to the sending server, then the identity token sent to the downstream server contains the certificate chain. With this information, the receiving server can perform its own certificate chain mapping and interoperability is increased with other vendors and platforms.

After the identity format is understood and parsed, the identity maps to a credential. For an ITTPincipal identity token, this identity maps one-to-one with the user ID fields.

For an ITTDistinguishedName identity token, the mapping depends on the user registry. For Lightweight Directory Access Protocol (LDAP), the configured search filter determines how the mapping occurs. For LocalOS, the first attribute of the distinguished name (DN), which is typically the same as the common name, maps to the user ID of the registry.

Identity assertion is only available using the Common Secure Interoperability Version 2 (CSIv2) protocol.

Related tasks

Chapter 5, “Authenticating users,” on page 111

The process of authenticating users involves a user registry and an authentication mechanism. Optionally, you can define trust between WebSphere Application Server and a proxy server, configure single sign-on capability, and specify how to propagate security attributes between application servers.

Identity assertions with trust validation

If you want an application or system provider to perform an identity assertion with trust validation, it can be accomplished by use of the Java Authentication and Authorization Service (JAAS) login framework, where trust validation is performed in one login module and credential creation in another. These two custom login modules are used to create a JAAS login configuration that performs a login to an identity assertion.

Two custom login module are required:

- A user-implemented trust association login module. This login module performs whatever trust verification the user requires. When trust is verified, the trust verification status and the login identity must be placed in a map in the share state of the login module to enable the credential creation login module to use that information. The map must be stored in the `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.state` property. State maps contain the following information:
 - `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.trusted` – set to `true`, if trusted, and `false`, if not trusted.
 - `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.principal` – contains the principal of the identity.
 - `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.certificates` – contains the certificate of the identity
- The `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule` module performs the credential creation. It requires that the trust state information be in the login context's shared state. This login module is protected by the Java 2 security runtime permissions for the following:
 - `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.initialize`
 - `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.login`

`IdentityAssertionLoginModule` searches for the trust information in the shared state property, `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.state`. This is a map that contains the trust status and the identity used to login. The map includes the following:

- `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.trusted` – if set to `true` it is trusted, `false` if not trusted.
- `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.principal` – if a principal is used, it contains the principal of the identity necessary to login.
- `com.ibm.wsspi.security.common.auth.module.IdentityAssertionLoginModule.certificates` – if a certificate is used, it contains an array of a certificate chain that includes the identity necessary to login.

A `WSLoginFailedException` is returned if the state, trust, or identity information is missing. The login module then performs a login of the identity. The subject now contains the new identity.

Message layer authentication

Defines the credential information and sends that information across the network so that a receiving server can interpret it.

When you send authentication information across the network using a token the transmission is considered message layer authentication because the data is sent with the message inside a service context.

A pure Java client uses Kerberos (KRB5) or basic authentication, or Generic Security Services Username Password (GSSUP), as the authentication mechanism to establish client identity.

However, a servlet can use either basic authentication (GSSUP) or the authentication mechanism of the server, Kerberos (KRB5) or Lightweight Third Party Authentication (LTPA), to send security information in the message layer. Use KRB5 or LTPA by authenticating or by mapping the basic authentication credentials to the security mechanism of the server.

The security token that is contained in a token-based credential is authentication mechanism-specific. The way that the token is interpreted is only known by the authentication mechanism. Therefore, each authentication mechanism has an object ID (OID) representing it. The OID and the client token are sent to the server, so that the server knows which mechanism to use when reading and validating the token. The following list contains the OIDs for each mechanism:

BasicAuth (GSSUP): oid:2.23.130.1.1.1
KRB5: OID: 1.2.840.113554.1.2.2
LTPA: oid:1.3.18.0.2.30.2
SWAM: No OID because it is not forwardable

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

On the server, the authentication mechanisms can interpret the token and create a credential, or they can authenticate basic authentication data from the client, and create a credential. Either way, the created credential is the *received* credential that the authorization check uses to determine if the user has access to invoke the method. You can specify the authentication mechanism by using the following property on the client side:

- com.ibm.CORBA.authenticationTarget

Basic authentication (BasicAuth) and KRB5 are currently the only valid values. You can configure the server through the administrative console.

Note: When **perform basic authentication** is enabled, if the client is not similarly configured (and does not pass a credential such as a user ID and password), the server object request broker (ORB) does not.

Configuring authentication retries

Situations occur where you want a prompt to display again if you entered your user ID and password incorrectly or you want a method to retry when a particular error occurs back at the client. If you can correct the error by information at the client side, the system automatically performs a retry without the client seeing the failure, if the system is configured appropriately.

Some of these errors include:

- Entering a user ID and password that are not valid
- Having an expired credential on the server
- Failing to find the stateful session on the server

By default, authentication retries are enabled and perform three retries before returning the error to the client. Use the com.ibm.CORBA.authenticationRetryEnabled property (True or False) to enable or disable authentication retries. Use the com.ibm.CORBA.authenticationRetryCount property to specify the number of retry attempts.

Immediate validating of a basic authentication login

In WebSphere Application Server Version 6.x, a behavior is defined during request_login for a BasicAuth login. In releases prior to Version 5, a BasicAuth login takes the user ID and password entered through the

loginSource method and creates a BasicAuth credential. If either the user ID or the password is not valid, the client program does not find out until the first method request is attempted. When the user ID or password is specified during a prompt or programmatic login, the user ID and password are authenticated by default with the security server, with a True or False returned as the result. If False, an org.omg.SecurityLevel2.LoginFailed exception is returned to the client indicating that the user ID and password are not valid. If True, then the BasicAuth credential is returned to the caller of the request_login. To disable this feature on the pure client, specify com.ibm.CORBA.validateBasicAuth=false. By default, this feature is set to True. On the server side, specify this property in the security dynamic properties.

Note: Set com.ibm.CORBA.validateBasicAuth=false whenever connecting to a z/OS server. This function does not currently work from a distributed client to a z/OS server since the SecurityServer is located using the "UNAUTHENTICATED" principal, which is not accepted on a z/OS system.

Chapter 6. Authorizing access to resources

WebSphere Application Server provides many different methods for authorizing accessing resources. For example, you can assign roles to users and configure a built-in or external authorization provider.

About this task

You can create an application, an Enterprise JavaBeans (EJB) module, or a Web module and secure them using assembly tools.

To authorize user or group access to resources, read the following articles:

1. Secure you application during assembly and deployment. For more information on how to create a secure application using an assembly tool, such as the IBM Rational Application Developer, see [Securing applications during assembly and deployment](#).
For general information about the tools that WebSphere Application Server supports, see [Assembly tools and Assembling applications](#).
2. Authorize access to Java Platform, Enterprise Edition (Java EE) resources. WebSphere Application Server supports authorization that is based on the Java Authorization Contract for Containers (JACC) specification in addition to the default authorization. When security is enabled in WebSphere Application Server, the default authorization is used unless a JACC provider is specified. For more information, see [“Authorization providers”](#) on page 522.
3. Authorize access to administrative resources. You can assign users and groups to predefined administrative roles such as the monitor, configurator, operator, administrator, auditor and iscadmins roles. These roles determine which tasks a user can perform in the administrative console. For more information, see [“Authorizing access to administrative roles”](#) on page 561.

What to do next

After authorizing access to resources, configure the Application Server for secure communication. For more information, see Chapter 7, [“Securing communications,”](#) on page 589.

Authorization technology

Authorization information determines whether a user or group has the necessary privileges to access resources.

WebSphere Application Server supports many authorization technologies including the following:

- Authorization involving the Web container and Java Platform, Enterprise Edition (Java EE) technology
- Authorization involving an enterprise bean application and Java EE technology
- Authorization involving Web services and Java EE technology
- Java Message Service (JMS)
- Java Authorization Contract for Containers (JACC)

WebSphere Application Server supports both a default authorization provider and an authorization provider that is based on the Java Authorization Contract for Containers (JACC) specification. The JACC-based authorization provider enables third-party security providers to handle the Java EE authorization. For more information, see [“JACC support in WebSphere Application Server”](#) on page 523.

- Java Authentication and Authorization Service (JAAS)
For more information, see [“Java Authentication and Authorization Service”](#) on page 371.
- Java 2 security
For more information, see [“Java 2 security”](#) on page 44.
- Naming and administrative authorization

- Pluggable authorization

WebSphere Application Server supports an authorization infrastructure that enables you to plug in an external authorization provider. For more information, see “Enabling an external JACC provider” on page 542.

Administrative roles and naming service authorization

WebSphere Application Server extends the Java Platform, Enterprise Edition (Java EE) security role-based access control to protect the product administrative and naming subsystems.

Administrative roles

A number of administrative roles are defined to provide the degrees of authority that are needed to perform certain WebSphere Application Server administrative functions from either the administrative console or the system management scripting interface called wsadmin. The authorization policy is only enforced when administrative security is enabled. The following table describes the administrative roles:

Table 17. Administrative roles that are available through the administrative console and wsadmin

Role	Description
Monitor	<p>An individual or group that uses the monitor role has the least amount of privileges. A monitor can complete the following tasks:</p> <ul style="list-style-type: none"> • View the WebSphere Application Server configuration. • View the current state of the Application Server.
Configurator	<p>An individual or group that uses the configurator role has the monitor privilege plus the ability to change the WebSphere Application Server configuration. The configurator can perform all the day-to-day configuration tasks. For example, a configurator can complete the following tasks:</p> <ul style="list-style-type: none"> • Create a resource. • Map an application server. • Install and uninstall an application. • Deploy an application. • Assign users and groups-to-role mapping for applications. • Set up Java 2 security permissions for applications. • Customize the Common Secure Interoperability Version 2 (CSIV2), Secure Authentication Service (SAS), and Secure Sockets Layer (SSL) configurations. <p>Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.</p>
Operator	<p>An individual or group that uses the operator role has monitor privileges plus ability to change the runtime state. For example, an operator can complete the following tasks:</p> <ul style="list-style-type: none"> • Stop and start the server. • Monitor the server status in the administrative console.

Table 17. Administrative roles that are available through the administrative console and wsadmin (continued)

Role	Description
Administrator	<p>An individual or group that uses the administrator role has the operator and configurator privileges plus additional privileges that are granted solely to the administrator role. For example, an administrator can complete the following tasks:</p> <ul style="list-style-type: none"> • Modify the server user ID and password. • Configure authentication and authorization mechanisms. • Enable or disable administrative security. <p>Note: In previous releases of WebSphere Application Server, the Enable administrative security option is known as the Enable global security option.</p> • Enforce Java 2 security using the Use Java 2 security to restrict application access to local resources option. • Change the Lightweight Third Party Authentication (LTPA) password and generate keys. • Create, update, or delete users in the federated repositories configuration. • Create, update, or delete groups in the federated repositories configuration. <p>Note: An administrator cannot map users and groups to the administrator roles.</p>
Adminsecuritymanager	<p>Only users who are granted this role can map users to administrative roles. Also, when fine-grained administrative security is used, only users who are granted this role can manage authorization groups. See “Administrative roles” on page 518 for more information.</p>
Deployer	<p>Users who are granted this role can perform both configuration actions and runtime operations on applications.</p>
Auditor	<p>Users granted this role can view and modify the configuration settings for the security auditing subsystem. For example, a user with the auditor role can complete the following tasks:</p> <ul style="list-style-type: none"> • Enable and disable the security auditing subsystem. • Select the event factory implementation to be used with the event factory plug-in point. • Select and configure the service provide, or emitter, or both to be used with the service provider plug-in point. • Set the audit policy that describes the behavior of the application server in the event of an error with the security auditing subsystem. • Define which security events are to be audited. <p>The auditor role includes the monitor role. This allows the auditor to view but not change the rest of the security configuration.</p>

Table 18. Additional administrative role that is available through the administrative console

Role	Description
iscadmins	<p>This role is only available for administrative console users and not for wsadmin users. Users who are granted this role have administrator privileges for managing users and groups in the federated repositories. For example, a user of the iscadmins role can complete the following tasks:</p> <ul style="list-style-type: none"> • Create, update, or delete users in the federated repositories configuration. • Create, update, or delete groups in the federated repositories configuration.

Table 19. Additional administrative role that is available through wsadmin

Role	Description
Deployer	<p>This role is only available for wsadmin users and not for administrative console users. Users who are granted this role can perform both configuration actions and run-time operations on applications.</p>

When administrative security is enabled, the administrative subsystem role-based access control is enforced. The administrative subsystem includes the security server, the administrative console, the wsadmin scripting tool, and all the Java Management Extensions (JMX) MBeans. When administrative security is enabled, both the administrative console and the administrative scripting tool require users to provide the required authentication data. Moreover, the administrative console is designed so the control functions that display on the pages are adjusted, according to the security roles that a user has. For example, a user who has only the monitor role can see only the non-sensitive configuration data. A user with the operator role can change the system state.

When you are changing registries (for example, from a federated repository to LDAP), make sure you remove the information that pertains to the previously configured registry for console users and console groups.

When administrative security is enabled, WebSphere Application Servers run under the server identity that is defined under the active user registry configuration. Although it is not shown on the administrative console and in other tools, a special Server subject is mapped to the administrator role. The WebSphere Application Server runtime code, which runs under the server identity, requires authorization to runtime operations. If no other user is assigned administrative roles, you can log into the administrative console or to the wsadmin scripting tool using the server identity to perform administrative operations and to assign other users or groups to administrative roles. Because the server identity is assigned to the administrative role by default, the administrative security policy requires the administrative role to perform the following operations:

- Change server ID and server password
- Enable or disable WebSphere Application Server administrative security
- Enforce Java 2 security using the **Use Java 2 security to restrict application access to local resources** option.
- Change the LTPA password or generate keys
- Assign users and groups to administrative roles

Primary administrative user name

The Version 6.1 release of WebSphere Application Server and subsequent releases require an administrative user, distinguished from the server user identity, to improve auditability of administrative actions. The user name specifies a user with administrative privileges that is defined in the local operating system.

Server user identity

The Version 6.1 release of WebSphere Application Server and subsequent releases distinguish the server identity from the administrative user identity to improve auditability. The server user identity is used for authenticating server-to-server communications.

Internal server ID

The internal server ID enables the automatic generation of the user identity for server-to-server authentication. Automatic generation of the server identity supports improved auditability for cells only for Version 6.1 or later nodes. In the Version 6.1 release of WebSphere Application Server, you can save the internally-generated server ID because the Security WebSphere Common Configuration Model (WCCM) model contains a new tag, internalServerId. You do not need to specify a server user ID and a password during security configuration except in a mixed-cell environment. An internally-generated server ID adds a further level of protection to the server environment because the server password is not exposed as it is in releases prior to Version 6.1. However, to maintain backwards compatibility, you must specify the server user ID if you use earlier versions of WebSphere Application Server.

When enabling security, you can assign one or more users and groups to naming roles. For more information, see Assigning users to naming roles. However, before assigning users to naming roles, configure the active user registry. User and group validation depends on the active user registry. For more information, see Configuring user registries.

Special subject

In addition to mapping users or groups, you can map a *special-subject* to the administrative roles. A special-subject is a generalization of a particular class of users. The AllAuthenticated or the AllAuthenticatedInTrustedRealms (when cross realm is involved) special subjects mean that the access check of the administrative role ensures that the user making the request is at least authenticated. The Everyone special subject means that anyone, authenticated or not, can perform the action as if security is not enabled.

Naming service authorization

CosNaming security offers increased granularity of security control over CosNaming functions. CosNaming functions are available on CosNaming servers such as the WebSphere Application Server. These functions affect the content of the WebSphere Application Server name space. Generally, you have two ways in which client programs result in CosNaming calls. The first is through the Java Naming and Directory Interface (JNDI) call. The second is with common object request broker architecture (CORBA) clients invoking CosNaming methods directly.

Four security roles are introduced :

- CosNamingRead
- CosNamingWrite
- CosNamingCreate
- CosNamingDelete

The roles have authority levels from low to high:

CosNamingRead

You can query the WebSphere Application Server name space, using, for example, the JNDI lookup method. The special-subject, Everyone, is the default policy for this role.

CosNamingWrite

You can perform write operations such as JNDI bind, rebind, or unbind, and CosNamingRead operations. As a default policy, Subjects are not assigned this role.

CosNamingCreate

You can create new objects in the name space through such operations as JNDI createSubcontext and CosNamingWrite operations. As a default policy, Subjects are not assigned this role.

CosNamingDelete

You can destroy objects in the name space, for example using the JNDI destroySubcontext method and CosNamingCreate operations. As a default policy, Subjects are not assigned this role.

A Server special-subject is assigned to all of the four CosNaming roles by default. The Server special-subject provides a WebSphere Application Server process, which runs under the server identity, to access all the CosNaming operations. The Server special-subject does not display and cannot be modified through the administrative console or other administrative tools.

Special configuration is not required to enable the server identity as specified when enabling administrative security for administrative use because the server identity is automatically mapped to the administrator role.

Users, groups, or the special subjects AllAuthenticated and Everyone can be added or removed to or from the naming roles from the WebSphere Application Server administrative console at any time. However, a server restart is required for the changes to take effect.

A best practice is to map groups or one of the special-subjects, rather than specific users, to naming roles because it is more flexible and easier to administer in the long run. By mapping a group to a naming role, adding or removing users to or from the group occurs outside of WebSphere Application Server and does not require a server restart for the change to take effect.

The CosNaming authorization policy is only enforced when administrative security is enabled. When administrative security is enabled, attempts to do CosNaming operations without the proper role assignment result in an org.omg.CORBA.NO_PERMISSION exception from the CosNaming server.

Each CosNaming function is assigned to only one role. Therefore, users who are assigned the CosNamingCreate role cannot query the name space unless they have also been assigned CosNamingRead. And in most cases a creator needs to be assigned three roles: CosNamingRead, CosNamingWrite, and CosNamingCreate. The CosNamingRead and CosNamingWrite roles assignment for the creator example are included in the CosNamingCreate role. In most of the cases, WebSphere Application Server administrators do not have to change the roles assignment for every user or group when they move to this release from a previous one.

Although the ability exists to greatly restrict access to the name space by changing the default policy, unexpected org.omg.CORBA.NO_PERMISSION exceptions can occur at runtime. Typically, Java EE applications access the name space and the identity they use is that of the user that authenticated to WebSphere Application Server when accessing the Java EE application. Unless the Java EE application provider clearly communicates the expected naming roles, use caution when changing the default naming authorization policy.

Administrative roles for business level applications

The Java 2 Platform, Enterprise Edition (J2EE) role-based authorization concept is extended to protect the WebSphere Application Server administrative subsystem. This protection applies to those administrative roles associated with business level applications.

Deploying business level applications on a server configured to hold business level applications requires a number of administrative roles that are defined to provide degrees of authority when performing certain administrative functions from either the Web-based administrative console or the system management scripting interface. The authorization policy is only enforced when administrative security is enabled. The following table describes the system management scripting command used for business level applications and the corresponding administrative role that is required in using the command:

Business level application - administrative roles

Command	Role Required
startBLA	Cell deployer, Cell operator, BLA deployer, BLA operator, Target deployer, Target operator
stopBLA	Cell deployer, Cell operator, BLA deployer, BLA operator, Target deployer, Target operator
createEmptyBLA	Cell configurator, Cell deployer
editBLA	Cell deployer, BLA deployer
viewBLA	Cell monitor, BLA monitor
listBLAs	Cell monitor, BLA monitor
deleteBLA	Cell configurator, Cell deployer, BLA developer
importAsset	Cell configurator, Cell deployer
editAsset	Cell configurator, Cell deployer, Asset deployer
viewAsset	Cell monitor, Asset monitor
listAsset	Cell monitor, Asset monitor
exportAsset	Cell monitor, Asset monitor

Business level application - administrative roles

deleteAsset	Cell configurator, Cell deployer, Asset deployer
updateAsset	Cell configurator, Cell deployer, Asset deployer
addCompUnit	Cell configurator Cell deployer, BLA deployer <i>(for the BLA to add the composition unit)</i> Asset-deployer <i>(for the asset to create the composition unit from)</i> Target-deployer <i>(for each target the composition unit is deployed to)</i> Relationship-deployer <i>(for each relationship the composition unit depends on that will result in creating a composition unit from the dependency asset)</i>
editCompUnit	Cell configurator, Cell deployer, BLA deployer <i>(for the BLA this composition unit belongs to)</i> Target deployer <i>(for each target that this composition unit is deployed to)</i>
viewCompUnit	Cell monitor, BLA monitor
listCompUnit	Cell monitor, BLA monitor
deleteCompUnit	Cell configurator, Cell deployer, BLA deployer <i>(for the BLA this composition unit belongs to)</i> Target deployer <i>(for each target that this composition unit is deployed to)</i>
setCompUnitTargetAutoStart	Cell configurator, Cell deployer
listControlOps	Cell monitor, BLA monitor
getBLAStatus	Cell monitor, BLA monitor
	<p>Where:</p> <ul style="list-style-type: none"> • BLA deployer specifies the deployer role for the BLA that is being managed. • BLA monitor specifies the monitor role for the BLA that is being managed. • BLA operator specifies the operator role for the BLA that is being managed. • Asset deployer specifies the deployer role for the asset that is being managed. • Asset monitor specifies the monitor role for the asset that is being managed. • Target deployer specifies the deployer for the target that the composition unit is being deployed to. • Target operator specifies the operator role for the target that the composition unit is being deployed to.

Related tasks

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.

Related reference

BLAManagement command group for the AdminTask object

You can use the Jython scripting language to configure and administer business-level applications with the wsadmin tool. Use the commands and parameters in the BLAManagement group to create, edit, export, delete, and query business-level applications in your configuration.

“Administrative roles” on page 518


The Java Platform, Enterprise Edition (Java EE) role-based authorization concept is extended to protect the WebSphere Application Server administrative subsystem.

Role-based authorization

Use authorization information to determine whether a caller has the necessary privileges to request a service.

The following figure illustrates the process that is used during authorization.

Authentication



Web resource access from a Web client is handled by a Web collaborator. The Enterprise JavaBeans (EJB) resource access from a Java client, whether an enterprise bean or a servlet, is handled by an EJB collaborator. The EJB collaborator and the Web collaborator extract the client credentials from the object request broker (ORB) current object. The client credentials are set during the authentication process as received credentials in the ORB current object. The resource and the received credentials are presented to the WSAccessManager access manager to check whether access is permitted to the client for accessing the requested resource.

The access manager module contains two main modules:

- The resource permission module helps determine the required roles for a given resource. This module uses a resource-to-roles mapping table that is built by the security runtime during application startup. To

build the resource-to-role mapping table, the security runtime reads the deployment descriptor of the enterprise beans or the Web module (`ejb-jar.xml` file or `web.xml` file)

- The authorization table module consults a role-to-user or group table to determine whether a client is granted one of the required roles. The role-to-user or group mapping table, also known as the *authorization table*, is created by the security runtime during application startup.

To build the authorization table, the security run time reads the application binding file, the `ibm-application-bnd.xmi` file, or the `ibm-application-bnd.xml` file, as appropriate.

Use authorization information to determine whether a caller has the necessary privilege to request a service. You can store authorization information many ways. For example, with each resource, you can store an access-control list, which contains a list of users and user privileges. Another way to store the information is to associate a list of resources and the corresponding privileges with each user. This list is called a *capability list*.

WebSphere Application Server uses the Java 2 Platform, Enterprise Edition (J2EE) authorization model. In this model, authorization information is organized as follows:

During the assembly of an application, permission to invoke methods is granted to one or more roles. A role is a set of permissions; for example, in a banking application, roles can include teller, supervisor, clerk, and other industry-related positions. The teller role is associated with permissions to run methods that are related to managing the money in an account, such as the withdraw and deposit methods. The teller role is not granted permission to close accounts; this permission is given to the supervisor role. The application assembler defines a list of method permissions for each role. This list is stored in the deployment descriptor for the application.

Three *special subjects* are not defined by the J2EE model: `AllAuthenticatedUsers`, `AllAuthenticatedInTrustedRealms` and `Everyone`. A special subject is a product-defined entity that is independent of the user registry. This entity is used to generically represent a class of users or groups in the registry.

- The `AllAuthenticatedUsers` subject permits all authenticated users to access protected methods. As long as the user can authenticate successfully, the user is permitted access to the protected resource. All of this is done independent of the user registry.
- The `AllAuthenticatedInTrustedRealms` subject permits all authenticated foreign users (those that are bound to other realms) to access protected methods. As long as the user can authenticate successfully, the user is permitted access to the protected resource. All of this is done independent of the user registry.
- `Everyone` is a special subject that permits unrestricted access to a protected resource. Users do not have to authenticate to get access; this special subject provides access to protected methods as if the resources are unprotected. All of this is done independent of the user registry.

During the deployment of an application, real users or groups of users are assigned to the roles. When a user is assigned to a role, the user gets all the method permissions that are granted to that role.

The application deployer does not need to understand the individual methods. By assigning roles to methods, the application assembler simplifies the job of the application deployer. Instead of working with a set of methods, the deployer works with the roles, which represent semantic groupings of the methods.

Users can be assigned to more than one role; the permissions that are granted to the user are the union of the permissions granted to each role. Additionally, if the authentication mechanism supports the grouping of users, these groups can be assigned to roles. Assigning a group to a role has the same effect as assigning each individual user to the role.

A best practice during deployment is to assign groups instead of individual users to roles for the following reasons:

- Improves performance during the authorization check. Typically far fewer groups exist than users.
- Provides greater flexibility, by using group membership to control resource access.

- Supports the addition and deletion of users from groups outside of the product environment. This action is preferred to adding and removing them to WebSphere Application Server roles. Stop and restart the enterprise application for these changes to take effect. This action can be very disruptive in a production environment.

At runtime, WebSphere Application Server authorizes incoming requests based on the user's identification information and the mapping of the user to roles. If the user belongs to any role that has permission to run a method, the request is authorized. If the user does not belong to any role that has permission, the request is denied.

The J2EE approach represents a declarative approach to authorization, but it also recognizes that you cannot deal with all situations declaratively. For these situations, methods are provided for determining user and role information programmatically. For enterprise beans, the following two methods are supported by WebSphere Application Server:

- **getCallerPrincipal:** This method retrieves the user identification information.
- **isCallerInRole:** This method checks the user identification information against a specific role.

For servlets, the following methods are supported by WebSphere Application Server:

- getRemoteUser
- isUserInRole
- getUserPrincipal

These methods correspond in purpose to the enterprise bean methods.

For more information on the J2EE security authorization model, see the following Web site:
<http://java.sun.com>

Administrative roles

The Java Platform, Enterprise Edition (Java EE) role-based authorization concept is extended to protect the WebSphere Application Server administrative subsystem.

A number of administrative roles are defined to provide degrees of authority that are needed to perform certain administrative functions from either the Web-based administrative console or the system management scripting interface. The authorization policy is only enforced when administrative security is enabled. The following table describes the administrative roles:

Administrative roles

Role	Description
Monitor	An individual or group that uses the monitor role has the least amount of privileges. A monitor can complete the following tasks: <ul style="list-style-type: none"> • View the WebSphere Application Server configuration. • View the current state of the Application Server.

Administrative roles

Configurator	<p>An individual or group that uses the configurator role has the monitor privilege plus the ability to change the WebSphere Application Server configuration. The configurator can perform all the daily configuration tasks. For example, a configurator can complete the following tasks:</p> <ul style="list-style-type: none"> • Create a resource. • Map an application server. • Install and uninstall an application. • Deploy an application. • Assign users and groups-to-role mapping for applications. • Set up Java 2 security permissions for applications. • Customize the Common Secure Interoperability Version 2 (CSIv2), Security Authentication Service (SAS), and Secure Sockets Layer (SSL) configurations. <p>Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.</p>
Operator	<p>An individual or group that uses the operator role has monitor privileges plus ability to change the runtime state. For example, an operator can complete the following tasks:</p> <ul style="list-style-type: none"> • Stop and start the server. • Monitor the server status in the administrative console.
Administrator	<p>An individual or group that uses the administrator role has the operator and configurator privileges, plus additional privileges that are granted solely to the administrator role. For example, an administrator can complete the following tasks:</p> <ul style="list-style-type: none"> • Modify the server user ID and password. • Configure authentication and authorization mechanisms. • Enable or disable administrative security. • Enable or disable Java 2 security. • Change the Lightweight Third Party Authentication (LTPA) password and generate keys. • Create, update, or delete users in the federated repositories configuration. • Create, update, or delete groups in the federated repositories configuration. <p>Note: An administrator cannot map users and groups to the administrator roles without also having the adminsecuritymanager role.</p>
iscadmins	<p>This role is only available for administrative console users, not for wsadmin users. Users who are granted this role have administrator privileges for managing users and groups in the federated repositories. For example, a user of the iscadmins role can complete the following tasks:</p> <ul style="list-style-type: none"> • Create, update, or delete users in the federated repositories configuration. • Create, update, or delete groups in the federated repositories configuration.
Deployer	<p>Users granted this role can complete both configuration actions and runtime operations on applications. See the “Deployer role” on page 520 section for more details.</p>
Admin Security Manager	<p>You can assign users and groups to the Admin Security Manager role on the cell level through wsadmin scripts and the administrative console. Using the Admin Security Manager role, you can assign users and groups to the administrative user roles and administrative group roles. However, an administrator cannot assign users and groups to the administrative user roles and administrative group roles including the Admin Security Manager role. See the “Admin Security Manager role” on page 521 section for more details.</p>

Administrative roles

Auditor	<p>Users granted this role can view and modify the configuration settings for the security auditing subsystem. For example, a user with the auditor role can complete the following tasks:</p> <ul style="list-style-type: none">• Enable and disable the security auditing subsystem.• Select the event factory implementation to be used with the event factory plug-in point.• Select and configure the service provide, or emitter. or both to be used with the service provider plug-in point.• Set the audit policy that describes the behavior of the application server in the event of an error with the security auditing subsystem.• Define which security events are to be audited. <p>The auditor role includes the monitor role. This allows the auditor to view but not change the rest of the security configuration. See the “Auditor role” on page 521 section for more details.</p>
---------	--

The server ID that is specified and the administrative ID, if specified, when enabling administrative security is automatically mapped to the administrator role.

Users and groups can be added or removed from administrative roles using the WebSphere Application Server administrative console by a user given the appropriate authority. The Primary administrative user name must be used to log on to the administrative console to change the administrative user and group roles other than the auditor role. Only a user with the auditor role can change the auditor user and group roles. When security auditing is initially enabled, the Primary administrative user is also given the auditor role, and can manage all of the administrative user and group roles including the those in the auditor role. A best practice is to map a group or groups, rather than specific users, to administrative roles because it is more flexible and easier to administer.

In addition to mapping user or groups, a special-subject can also be mapped to the administrative roles. A special-subject subject is a generalization of a particular class of users. The AllAuthenticated special subject means that the access check of the administrative role ensures that the user making the request is at least authenticated. The Everyone special subject means that anyone, authenticated or not, can perform the action, as if security was not enabled.

Deployer role

A user that is granted a deployer role can complete all of the configuration and runtime operations on an application. A deployer role can be subsets of both configurator and operator roles. However, a user granted a deployer role cannot configure or operate any other resources, such as a server, node.

When fine-grained administrative security is used, only a user granted a deployer role to an application can configure and operate that application.

Cell-level configurators can configure applications. Cell-level operators can also operate (start and stop) applications. However, a user granted a deployer role at cell level can also complete configuration and operation on all applications.

The following table lists the capabilities of the deployer role when fine-grained administrative security is used:

Operation	Required Roles (Any one)
Install application	Cell-configurator, application-deployer, target-deployer
Uninstall application	Cell-configurator, application-deployer, target-deployer

Operation	Required Roles (Any one)
List application	Cell-monitor, application-monitor
Edit, update and redeploy application	Cell-configurator, application-deployer
Export application	Cell-monitor, application-monitor
Start or stop application	Cell-operator, application-deployer

Where:

Cell-configurator

Specifies the configurator role at cell level.

Application-deployer

Specifies the deployer role for the application that is being managed.

Target-deployer

Specifies the deployer role for all servers or clusters for which an application is targeted. If you have a target-deployer role, you can install a new application on the target. However, to edit or update the installed application, you must be included in the authorization group of the installed application-deployer.

The target-deployer cannot explicitly start or stop a new application. However, when a target-deployer starts a server on a target, all of the applications that have their auto-start attribute set to yes are started when the server starts.

It is recommended that the application-deployer set this attribute to true if the application-deployer does not want the application to be started by the target-deployer.

Admin Security Manager role

The Admin Security Manager role separates administrative security administration from other application administration.

By default, serverId and adminID, if specified, are assigned to this role in the cell level authorization table. This role implies a monitor role. However, an administrator role does not imply the Admin Security Manager role.

When fine-grained admin security is used, only a user granted this role at cell level can manage administrative authorization groups. However, a user granted this role for each administrative authorization group can map users to administrative roles for those groups. The following list summarizes the capabilities of the Admin Security Manager role at different levels, such as the cell and administrative authorization group levels.

Action	Role
Map users to administrative roles for cell level	Only the Admin Security Manager of the cell
Map users to administrative roles for an authorization group	Only the Admin Security Manager of that authorization group or the Admin Security Manager of the cell
Manage authorization groups, create, delete, add resource to an authorization group, or remove resource from an authorization group or list	Only the Admin Security Manager of the cell

Auditor role

The auditor role separates security auditing administration from administrative security and other application administration.

The auditor role was added to allow distinct separation of the authority of an auditor from the authority of the administrator. The auditor role can be granted to administrators to combine their authority. When security is first enabled, the auditor role is assigned to the primary administrator. If in your situation the separation of authority is required, administrators can remove the auditor role from themselves and assign the auditor role to other users.

A fine grained security for the auditor role is not implemented, which results in the auditor role requiring the monitor role. This process allows the auditor to read but not modify the panels managed by the administrator. The auditor has full authority to read and modify the panels associated with the security auditing subsystem. The administrator will have the monitor role for those panels, however, the administrator cannot modify those panels.

Related tasks

“Assigning users to naming roles” on page 565

Use this task to assign users to naming roles by using the administrative console.

Authorization providers

WebSphere Application Server supports authorization that is based on the Java Authorization Contract for Containers (JACC) specification in addition to the default authorization.

JACC is a specification introduced in Java Platform, Enterprise Edition (Java EE) 1.4. It enables third-party security providers to manage authorization in the application server.

Note: In WebSphere Application Server version 7.0, Java Authorization Contract for Containers (JACC) specification 1.4 is being applied. In JACC specification 1.4, we support Java EE5 that includes Servlet 2.5 and EJB 3. The biggest functional change with the introduction of JACC specification 1.4 is the inclusion of annotations for propagating security policy information.

When security is enabled in WebSphere Application Server, the default authorization is used unless a JACC provider is specified. The default authorization does not require special setup, and the default authorization engine makes all of the authorization decisions. However, if a JACC provider is configured and set up for WebSphere Application Server to use, all of the enterprise beans and Web authorization decisions are delegated to the JACC provider.

WebSphere Application Server supports security for Java EE applications and also for its administrative components. Java EE applications, such as Web and Enterprise JavaBeans (EJB) components are protected and authorized per the Java EE specification. The administrative components are internal to WebSphere Application Server and are protected by the role-based authorizer. The administrative components include the administrative console, MBeans, and other components such as naming and security. For more information on administrative security, see “Role-based authorization” on page 516.

When a JACC provider is used for authorization in WebSphere Application Server, all of the Java EE application-based authorization decisions are delegated to the provider per the JACC specification. However, all administrative security authorization decisions are made by the WebSphere Application Server default authorization engine. The JACC provider is not called to make the authorization decisions for administrative security.

When a protected Java EE resource is accessed, the authorization decision to give access to the principal is the same whether using the default authorization engine or a JACC provider. Both of the authorization models satisfy the J2EE specification, and function the same. Choose a JACC provider only when you want to work with an external security provider such as Tivoli Access Manager. In this instance, the security provider must support the JACC specification and be set up to work with WebSphere Application Server. Setting up and configuring a JACC provider requires additional configuration steps, depending on the provider. Unless you have an external security provider that you can use with WebSphere Application Server, use the default authorization.

JACC support in WebSphere Application Server

WebSphere Application Server supports the Java Authorization Contract for Containers (JACC) specification, which enables third-party security providers to handle the Java Platform, Enterprise Edition (Java EE) authorization.

The JACC specification requires that both the containers in the application server and the provider satisfy some requirements. Specifically, the containers are required to propagate the security policy information to the provider during the application deployment and to call the provider for all authorization decisions. The providers are required to store the policy information in their repository during application deployment. The providers then use this information to make authorization decisions when called by the container.

JACC access decisions

When security is enabled and an enterprise bean or Web resource is accessed, the Enterprise JavaBeans (EJB) container or Web container calls the security runtime to make an authorization decision on whether to permit access. When using an external provider, the access decision is delegated to that provider.

According to the Java Authorization Contract for Containers (JACC) specification, the appropriate permission object is created, the appropriate policy context handlers are registered, and the appropriate policy context identifier (contextID) is set. A call is made to the `java.security.Policy` object method that is implemented by the provider to make the access decision.

The following sections describe how the provider is called for both the enterprise bean and the Web resources.

Access decisions for enterprise beans

When security is enabled, and an EJB method is accessed, the EJB container delegates the authorization check to the security runtime. If JACC is enabled, the security runtime uses the following process to perform the authorization check:

1. Creates the `EJBMethodPermission` object using the bean name, method name, interface name, and the method signature.
2. Creates the context ID and sets it on the thread by using the `PolicyContext.setContextID(contextID)` method.
3. Registers the required policy context handlers, including the Subject policy context handler.
4. Creates the `ProtectionDomain` object with principal in the Subject. If no principal exists, null is passed for the principal name.
5. The access decision is delegated to the JACC provider by calling the `implies` method of the `Policy` object, which is implemented by the provider. The `EJBMethodPermission` and the `ProtectionDomain` objects are passed to this method.
6. The `isCallerInRole` access check also follows the same process, except that an `EJBRoleRefPermission` object is created instead of an `EJBMethodPermission` object.

Access decisions for Web resources

When security is enabled and configured to use a JACC provider, and when a Web resource such as a servlet or a JavaServer Pages (JSP) file is accessed, the security runtime delegates the authorization decision to the JACC provider by using the following process:

1. A `WebResourcePermission` object is created to see if the URI is cleared. If the provider honors the Everyone subject it is also selected here.
 - a. The `WebResourcePermission` object is constructed with the `urlPattern` and the HTTP method accessed.
 - b. A `ProtectionDomain` object with a null principal name is created.

- c. The JACC provider `Policy.implies` method is called with the permission and the protection domain. If the URI access is cleared or given access to the Everyone subject, the provider permits access (return true) in the `implies` method. Access is then granted without further checks.
2. If access is not granted in the previous step, a `WebUserDataPermission` object is created and used to see if the Uniform Resource Identifier (URI) is precluded, excluded or must be redirected using the HTTPS protocol.
 - a. The `WebUserDataPermission` object is constructed with the `urlPattern` accessed, the HTTP method invoked, and the transport type of the request. If the request is over HTTPS, the transport type is set to `CONFIDENTIAL`; otherwise, null is passed.
 - b. A `ProtectionDomain` object with a null principal name is created.
 - c. The JACC provider `Policy.implies` method is called with the permission and the protection domain. If the request is using the HTTPS protocol and the `implies` method returns `false`, the HTTP 403 error is returned to imply excluded and precluded permission. In this case no further checks are performed. If the request is not using the HTTPS protocol, and the `implies` returns `false`, the request is redirected over HTTPS.
3. The security runtime attempts to authenticate the user. If the authentication information already exists (for example, LTPA token), it is used. Otherwise, the user's credentials must be entered.
4. After the user credentials are validated, a final authorization check is performed to see if the user is granted access privileges to the URI.
 - a. As in Step 1, the `WebResourcePermission` object is created. The `ProtectionDomain` object now contains the Principal that is attempting to access the URI. The Subject policy context handler also contains the user's information, which can be used for the access check.
 - b. The provider `implies` method is called using the `Permission` object and the `ProtectionDomain` object created previously. If the user is granted permission to access the resource, the `implies` method returns `true`. If the user is not granted access, the `implies` method returns `false`.

Even if the order listed previously is changed later (for example, to improve performance) the end result is the same. For example, if the resource is precluded or excluded, the end result is that the resource cannot be accessed.

Using information from the Subject for access decision

If the provider relies on the WebSphere Application Server generated Subject for access decision, the provider can query the public credentials in the Subject to obtain the `WSCredential` credential. The `WSCredential` API is used to obtain information about the user, including the name and the groups that the user belongs to. This information is used to make the access decision.

If the provider adds the required information to the Subject, WebSphere Application Server can use the information to make the access decision. The provider might add the information by using the Trust Association Interface feature or by plugging login modules into the Application Server.

The security attribute propagation section contains additional documentation on how to add the WebSphere Application Server required information to the Subject. For more information, see "Propagating security attributes among application servers" on page 453.

Dynamic module updates in JACC

WebSphere Application Server supports dynamic updates to Web modules under certain conditions. If a Web module is updated, deleted or added to an application, only that module is stopped and started as appropriate. The other existing modules in the application are not impacted, and the application itself is not stopped and then restarted.

When using the default authorization engine, any security policies are modified in the Web modules and the application is stopped and then restarted. When using the Java Authorization Contract for Containers

(JACC) based authorization, the behavior depends on the functionality that a provider supports. If a provider can handle dynamic changes to the Web modules, then only the Web modules are impacted. Otherwise, the entire application is stopped and restarted for the new changes in the Web modules to take effect.

A provider can indicate if it supports the dynamic updates by configuring the **Supports dynamic module updates** option in the JACC configuration model (see “Authorizing access to J2EE resources using Tivoli Access Manager” on page 537 for more information). This option can be enabled or disabled using the administrative console or by scripting. It is expected that most providers store the policy information in their external repository, which makes it possible for them to support these dynamic updates. This option should be enabled by default for most providers.

When the **Supports dynamic module updates** option is enabled, if a Web module that contains security roles is dynamically added, modified, or deleted, only the specific Web modules are impacted and restarted. If the option is disabled, the entire application is restarted. When dynamic updates are performed, the security policy information of the modules impacted are propagated to the provider. For more information about security policy propagation, see “JACC policy propagation” on page 527.

Initialization of the JACC provider

If a Java Authorization Contract for Containers (JACC) provider requires initialization during server startup, for example, to enable the client code to communicate to the server code, the provider can implement the `com.ibm.wsspi.security.authorization.InitializeJACCProvider` interface. See “Interfaces that support JACC” on page 555 for more information.

When this interface is implemented, it is called during server startup. Any custom properties in the JACC configuration model are propagated to the `initialize` method of this implementation. The custom properties can be entered using either the administrative console or by scripting.

During server shutdown, the `cleanup` method is called for any clean-up work that a provider requires. Implementation of this interface is strictly optional, and is used only if the provider requires initialization during server startup.

Mixed node environment and JACC

Authorization using Java Authorization Contract for Containers (JACC) is a new feature in WebSphere Application Server Version 6.0.x. Also, the JACC configuration is set up at the cell level and is applicable for all the nodes and servers in that cell.

If you are planning to use the JACC-based authorization, the cell must contain Version 6.0.x and later nodes only. This restriction implies that a mixed node environment containing a set of Version 5.x nodes in a Version 6.0.x or later cell is not supported.

JACC providers

The Java Authorization Contract for Containers (JACC) is a specification that was first introduced in Java Platform, Enterprise Edition (Java EE) Version 1.4 through the Java Specifications Request (JSR) 115 process. JACC specification 1.4 is included for WebSphere Application Server version 7.0 for Java EE 5 support.. This specification defines a contract between Java EE 5 containers and authorization providers.

The contract enables third-party authorization providers to plug into Java EE 5 application servers, such as WebSphere Application Server, to make the authorization decisions when a Java EE 5 resource is accessed. The access decisions are made through the standard `java.security.Policy` object.

To plug in to WebSphere Application Server, the third-party JACC provider must implement the policy class, policy configuration factory class, and policy configuration interface, which are all required by the JACC specification.

The JACC specification does not specify how to handle the authorization table information between the container and the provider. It is the responsibility of the provider to provide some management facilities to handle this information. The container is not required to provide the authorization table information in the binding file to the provider.

WebSphere Application Server provides the RoleConfigurationFactory and the RoleConfiguration role configuration interfaces to help the provider obtain information from the binding file, as well as an initialization interface (InitializeJACCProvider). The implementation of these interfaces is optional. See “Interfaces that support JACC” on page 555 for more information about these interfaces.

Tivoli Access Manager as the default JACC provider for WebSphere Application Server

The JACC provider in WebSphere Application Server is implemented by both the client and the server pieces of the Tivoli Access Manager. The client piece of Tivoli Access Manager is embedded in WebSphere Application Server. The server piece is located on a separate installable CD that is shipped as part of the WebSphere Application Server Network Deployment (ND) package.

The JACC provider is not the default authorization. You must configure WebSphere Application Server to use the JACC provider.

JACC policy context handlers

WebSphere Application Server supports all of the policy context handlers that are required by the Java Authorization Contract for Containers (JACC) specification. However, due to performance impacts, the Enterprise JavaBeans (EJB) arguments policy context handler is not activated unless it is specifically required by the provider. Performance impacts result if objects must be created for each argument of each EJB method.

If the provider supports and requires this context handler, select the **Requires the EJB arguments policy context handler for access decisions** check box in the External JACC provider link under the Authorization providers panel or by using scripting. Any changes to this option are effective after the servers are restarted. By default this option is disabled. Disable this option when using Tivoli Access Manager as the JACC provider, because the argument values are not required for access decisions.

JACC policy context identifiers (ContextID) format

A policy context identifier is defined as a unique string that represents a policy context. A policy context contains all of the security policy statements as defined by the Java Contract for Containers (JACC) specification that affect access to the resources in a Web or Enterprise JavaBeans (EJB) module.

During policy propagation to the JACC provider, a PolicyConfiguration object is created for each policy context. The object is populated with the policy statements, represented by the JACC permission objects that correspond to the context. The object is propagated to the JACC provider using the JACC specification APIs.

WebSphere Application Server makes the contextID unique by using the href:cellName/appName/moduleName string as the contextID format for the modules. The href part of the string indicates that a hierarchical name is passed as the context ID. The cellName represents the name of the deployment manager cell or the base cell where the application is installed. After an application is installed in one cell (for example, in a base application server where the cell name is base1) and is added to a deployment manager cell whose name is cell1 by using addNode, the context ID for the modules in the application contain base1 (not cell1) as the cell name because the application is initially installed in base1.

The appName part of the string in the context ID represents the application name containing the module. The moduleName refers to the name of the module.

As an example, the context ID for the module Increment.jar file in an application named DefaultApplication that is installed in cell1 is the href:cell1/DefaultApplication/Increment.jar file.

JACC policy propagation

When an application is installed or deployed in WebSphere Application Server, the security policy information in the application is propagated to the provider when the configuration is saved. The context ID for the application is saved in its `application.xml` file, that is used for propagating the policy to the Java Authorization Contract for Containers (JACC) provider, and also for access decisions for Java Platform, Enterprise Edition (Java EE) resources.

When an application is uninstalled, the security policy information in the application is removed from the provider when the configuration is saved.

If the provider implemented the `RoleConfiguration` interface, the security policy information that is propagated to the policy provider also contains the authorization table information. See “Interfaces that support JACC” on page 555 for more information about this interface.

If an application does not contain security policy information, the `PolicyConfiguration` (and the `RoleConfiguration`, if implemented) objects do not contain any information. The existence of empty `PolicyConfiguration` and `RoleConfiguration` objects indicates that security policy information for the module does not exist.

After an application is installed, it can be updated without being uninstalled and reinstalled. For example, a new module can be added to an existing application, or an existing module can be modified. In this instance, the information in the impacted modules is propagated to the provider by default. A module is impacted when the deployment descriptor of the module or annotations within the module are changed as part of the update. If the provider supports the `RoleConfiguration` interfaces, the entire authorization table for that application is propagated to the provider.

If the security information is not propagated to the provider during application updates, you can set the `com.ibm.websphere.security.jacc.propagateonappupdate` Java virtual machine (JVM) property to `false` in the deployment manager, in a Network Deployment environment, or the unmanaged base application server. If this property is set to `false`, any updates to an existing application in the server are not propagated to the provider. You also can set this property on a per-application basis using the custom properties of an application. The `wsadmin` tool can be used to set the custom property of an application. If this property is set at the application level, any updates to that application are not propagated to the provider. If the update to an application is a full update, for example, a new application enterprise archive (EAR) file is used to replace the existing one, and the provider is refreshed with the entire application security policy information.

When an application is installed and saved, the security policy information in that application is updated in the provider from the deployment manager. The application is not propagated to its respective nodes until the synchronization command is issued and completed. Also, when an application is uninstalled and saved at the deployment manager, the policy for that application is removed from the JACC provider.

However, unless the synchronization command is issued and completed from the deployment manager to the nodes hosting the application, the applications are still running in the respective nodes. In this instance, deny any access to this application because the JACC provider does not contain the required information to make the access decision for that application. Any updates to the application already installed as described previously are also propagated to the provider from the deployment manager. The changes in the provider are not synchronized with the applications in the nodes until the synchronization is completed.

As mentioned earlier, the security policy information is propagated to the JACC provider during the save operation. The `SystemOut.log` file indicates the success or failure of the propagation to the provider. Check the log file after the installation to ensure that the propagation had no problems. If the propagation had any problems, access to the application fails when Tivoli Access Manager is used as the JACC provider.

If the security policy information for the application is successfully propagated to the provider, the audit statements with the message key SECJ0415I appear. However, if there was a problem propagating the security policy information to the provider (for example: network problems, JACC provider is not available), the SystemOut.log files contain the error message with the message keys SECJ0396E during install or SECJ0398E during modification. The installation of the application is not stopped due to a failure to propagate the security policy to the JACC provider. Also, in the case of failure, no exception or error messages appear during the save operation. When the problem causing this failure is fixed, run the **propagatePolicyToJaccProvider** tool to propagate the security policy information to the provider without reinstalling the application. For more information, see “Propagating security policy of installed applications to a JACC provider using wsadmin scripting” on page 1054.

JACC registration of the provider implementation classes

The JACC specification states that providers can plug in their provider using the `javax.security.jacc.policy.provider` and the `javax.security.jacc.PolicyConfigurationFactory.provider` system properties.

The `javax.security.jacc.policy.provider` property is used to set the policy object of the provider, while the `javax.security.jacc.PolicyConfigurationFactory.provider` property is used to set the provider `PolicyConfigurationFactory` implementation.

Although both system properties are supported in WebSphere Application Server, it is highly recommended that you use the configuration model that is provided. You can set these values using either the JACC configuration panel (see “Authorizing access to J2EE resources using Tivoli Access Manager” on page 537 for more information) or by using wsadmin scripting. One of the advantages of using the configuration model instead of the system properties is that the information is entered in one place at the cell level, and is propagated to all nodes during synchronization. Also, as part of the configuration model, additional properties can be entered, as described in the JACC configuration panel.

Using the configuration model is especially recommended in the case of a Network Deployment (ND) environment where multiple application servers can exist in the configuration. If the system properties are used, you must ensure that each of the Java virtual machine (JVM) processes in the configuration set these properties. If the configuration model is used, the information is propagated to all processes through the synchronization process of the application server.

Role-based security with embedded Tivoli Access Manager

The Java Platform, Enterprise Edition (Java EE) role-based authorization model uses the concepts of roles and resources. An example is provided here.

Roles	Methods		
	getBalance	deposit	closeAccount
Teller	granted	granted	
Cashier	granted		
Supervisor			granted

In the example of the banking application that is conceptualized in the previous table, three roles are defined: teller, cashier, and supervisor. Permission to perform the `getBalance`, `deposit`, and `closeAccount` application methods are mapped to these roles. From the example, you can see that users assigned the role, Supervisor, can run the `closeAccount` method, whereas the other two roles are unable to run this method.

The term, principal, within WebSphere Application Server security refers to a person or a process that performs activities. *Groups* are logical collections of principals that are configured in WebSphere

Application Server to promote the ease of applying security. Roles can be mapped to principals, groups, or both. The entry that is invoked in the following table indicates that the principal or group can invoke any methods that are granted to that role.

Principal/Group	Roles		
	Teller	Cashier	Supervisor
TellerGroup	Invoke		
CashierGroup		Invoke	
SupervisorGroup			
Frank: A principal who is not a member of any of the previous groups		Invoke	Invoke

In the previous example, the principal Frank, can invoke the `getBalance` and the `closeAccount` methods, but cannot invoke the `deposit` method because this method is not granted either the Cashier or the Supervisor role.

At the time of application deployment, the Java Authorization Contract for Container (JACC) provider of Tivoli Access Manager populates the Tivoli Access Manager-protected object space with any security policy information that is contained in the application deployment descriptor and or annotations. This security information is used to determine access whenever the WebSphere Application Server resource is requested.

By default, the Tivoli Access Manager access check is performed using the role name, the cell name, the application name, and the module name.

Tivoli Access Manager access control lists (ACLs) determine which application roles are assigned to a principal. ACLs are attached to the applications in the Tivoli Access Manager-protected object space at the time of application deployment.

Principal-to-role mappings are managed from the WebSphere Application Server administrative console and are never modified using Tivoli Access Manager. Direct updates to ACLs are performed for administrative security users only.

The following sequence of events occur:

1. During application deployment, policy information is sent to the JACC provider of Tivoli Access Manager . This policy information contains permission-to-role mappings and role-to-principal and role-to-group mapping information.
2. The JACC provider of Tivoli Access Manager converts the information into the required format, and passes this information to the Tivoli Access Manager policy server.
3. The policy server adds entries to the Tivoli Access Manager-protected object space to represent the roles that are defined for the application and the permission-to-role mappings. A permission is represented as a Tivoli Access Manager-protected object and the role that is granted to this object is attached as an extended attribute.

Tivoli Access Manager integration as the JACC provider

Tivoli Access Manager uses the Java Authorization Contract for Container (JACC) model in WebSphere Application Server to perform access checks.

Tivoli Access Manager consists of the following components:

- Run time
- Client configuration
- Authorization table support

- Access check
- Authentication using the PDLoginModule module

Tivoli Access Manager run-time changes that are used to support JACC

For the run-time changes, Tivoli Access Manager implements the PolicyConfigurationFactory and the PolicyConfiguration interfaces, as required by JACC. During the application installation, the security policy information in the deployment descriptor and the authorization table information in the binding files are propagated to the Tivoli provider using these interfaces. The Tivoli provider stores the policy and the authorization table information in the Tivoli Access Manager policy server by calling the respective Tivoli Access Manager application programming interfaces (API).

Tivoli Access Manager also implements the RoleConfigurationFactory and the RoleConfiguration interfaces. These interfaces are used to ensure that the authorization table information is passed to the provider with the policy information. See “Interfaces that support JACC” on page 555 for more information about these interfaces.

Tivoli Access Manager client configuration

To configure the Tivoli Access Manager client, you can use either the administrative console or wsadmin scripting. You can access the administrative console panels for the Tivoli Access Manager client configuration by clicking **Security > Global security > External authorization providers**. Under Related Items, click **External JACC provider**. The Tivoli client must be set up to use the Tivoli Access Manager JACC Provider.

For more information about how to configure the Tivoli Access Manager client, see “Tivoli Access Manager JACC provider configuration” on page 545.

Authorization table support

Tivoli Access Manager uses the RoleConfiguration interface to ensure that the authorization table information is passed to the Tivoli Access Manager provider when the application is installed or deployed. When an application is deployed or edited, the set of users and groups for the user or group-to-role mapping are obtained from the Tivoli Access Manager server, which shares the same Lightweight Directory Access Protocol (LDAP) server as WebSphere Application Server. This sharing is accomplished by plugging into the application management users or groups-to-role administrative console panels. The management APIs are called to obtain users and groups rather than relying on the WebSphere Application Server-configured LDAP registry.

The user or group-to-role mapping is on the application level, not on the node level.

Access check

When WebSphere Application Server is configured to use the JACC provider for Tivoli Access Manager, it passes the information to Tivoli Access Manager to make the access decision. The Tivoli Access Manager policy implementation queries the local replica of the access control list (ACL) database for the access decision.

Authentication using the PDLoginModule module

The custom login module in WebSphere Application Server can do the authentication. This login module is plugged in before the WebSphere Application Server-provided login modules. The custom login modules can provide information that can be stored in the Subject. If the required information is stored, no additional registry calls are made to obtain that information.

As part of the JACC integration, the Tivoli Access Manager-provided PDLoginModule module is also used to plug into WebSphere Application Server for Lightweight Third Party Authentication (LTPA), Kerberos (KRB5) and Simple WebSphere Authentication Mechanism (SWAM) authentication. The PDLoginModule module is modified to authenticate with the user ID or password. The module is also used to fill in the required attributes in the Subject so that no registry calls are made by the login modules in WebSphere Application Server. The information that is placed in the Subject is available for the Tivoli Access Manager policy object to use for access checking.

Note: SWAM is deprecated in WebSphere Application Server Version 7.0 and will be removed in a future release.

Note: When using Kerberos authentication mechanism and Tivoli Access Manager, TAM loginModule creates the PDPrincipal without first going through the Tivoli Access Manager authentication process. Also when using Kerberos authentication mechanism and Tivoli Access Manager, the Tivoli Access Manager policy is not enforced in WebSphere Application Server Version 7.0.

Tivoli Access Manager security for WebSphere Application Server

WebSphere Application Server provides embedded IBM Tivoli Access Manager client technology to secure your WebSphere Application Server-managed resources.

The benefits of using Tivoli Access Manager that are described here are only applicable when Tivoli Access Manager client code is used with the Tivoli Access Manager server:

- Robust container-based authorization
- Centralized policy management
- Management of common identities, user profiles, and authorization mechanisms
- Single-point security management for Java Platform, Enterprise Edition (Java EE) compliant and non-compliant Java EE resources using the administrative console for Tivoli Access Manager Web Portal Manager
- No requirements for coding or deployment changes to applications
- Easy management of users, groups, and roles using the WebSphere Application Server administrative console


WebSphere Application Server supports the Java Authorization Contract for Containers (JACC) specification. JACC details the contract requirements for Java EE containers and authorization providers. With this contract, authorization providers can perform the access decisions for resources in Java EE application servers such as WebSphere Application Server. The Tivoli Access Manager security utility that is embedded within WebSphere Application Server is JACC-compliant and is used to:

- Add security policy information when applications are deployed
- Authorize access to WebSphere Application Server-secured resources.

When applications are deployed, the embedded Tivoli Access Manager client takes any policy and or user and role information that is stored within the application deployment descriptor or using annotations and stores it within the Tivoli Access Manager Policy Server.

The Tivoli Access Manager JACC provider is also called when a user requests access to a resource that is managed by WebSphere Application Server.

Embedded Tivoli Access Manager client architecture



The previous figure illustrates the following sequence of events:

1. Users that access protected resources are authenticated using the Tivoli Access Manager login module that is configured for use when the embedded Tivoli Access Manager client is enabled.
2. The WebSphere Application Server container uses information from the Java EE application deployment descriptor and annotations to determine the required role membership.
3. WebSphere Application Server uses the embedded Tivoli Access Manager client to request an authorization decision from the Tivoli Access Manager authorization server. Additional context information, when present, is also passed to the authorization server. This context information is comprised of the cell name, Java EE application name, and Java EE module name. If the Tivoli Access Manager policy database has policies that are specified for any of the context information, the authorization server uses this information to make the authorization decision.
4. The authorization server consults the permissions that are defined for the specified user within the Tivoli Access Manager-protected object space. The protected object space is part of the policy database.
5. The Tivoli Access Manager authorization server returns the access decision to the embedded Tivoli Access Manager client.
6. WebSphere Application Server either grants or denies access to the protected method or resource, based on the decision that is returned from the Tivoli Access Manager authorization server.

At its core, Tivoli Access Manager provides an authentication and authorization framework. You can learn more about Tivoli Access Manager, including the information that is necessary to make deployment decisions, by reviewing the product documentation. The following guides are available at the IBM Tivoli Access Manager for e-business information center:

- *IBM Tivoli Access Manager Base Installation Guide*

This guide describes how to plan, install, and configure a Tivoli Access Manager secure domain. Using a series of easy installation scripts, you can quickly deploy a fully functional secure domain. These scripts are very useful when prototyping the deployment of a secure domain.

To access this guide in the IBM Tivoli Access Manager for e-business information center, click **Access Manager for e-business > Base Information > Base Installation Guide**.

- *IBM Tivoli Access Manager Base Administration Guide*


This document presents an overview of the Tivoli Access Manager security model for managing protected resources. This guide describes how to configure the Tivoli Access Manager servers that

make access control decisions. In addition, detailed instructions describe how to perform important tasks, such as declaring security policies, defining protected object spaces, and administering user and group profiles.

To access this guide in the IBM Tivoli Access Manager for e-business information center, click **Access Manager for e-business > Base Information > Base Administration Guide**.

Tivoli Access Manager provides centralized administration of multiple servers.

WebSphere Application Server Cell



The previous figure is an example architecture showing WebSphere Application Servers secured by Tivoli Access Manager.

The participating WebSphere Application Servers use a local replica of the Tivoli Access Manager policy database to make authorization decisions for incoming requests. The local policy databases are replicas of the master policy database. The master policy database is installed as part of the Tivoli Access Manager installation. Having policy database replicas on each participating WebSphere Application Server node optimizes performance when making authorization decisions and provides failover capability.

Although the authorization server can also be installed on the same system as WebSphere Application Server, this configuration is not illustrated in the diagram.

All instances of Tivoli Access Manager and WebSphere Application Server in the example architecture share the Lightweight Directory Access Protocol (LDAP) user registry on Machine E.

The LDAP registries that are supported by WebSphere Application Server are also supported by Tivoli Access Manager.

It is possible to have separate WebSphere Application Server profiles on the same host that is configured for different Tivoli Access Manager servers. Such an architecture requires that the profiles are configured for separate Java SE Runtime Environments (JRE 6) and therefore you need multiple JREs installed on the same host.

Security Annotations

Annotations are a powerful programming resulting from the JSR-175 recommendation. An annotation is a standard way to include supported security behaviors while allowing, the source code and configuration files to be generated automatically.

In Java Platform, Enterprise Edition (Java EE) 5, The security roles and policies can be defined using annotations as well as within the deployment descriptor. During the installation of the application, the security policies and roles defined using annotations are merged with the security policies and roles defined within the deployment descriptor. This merge is performed by the Annotations Metadata Manager (AMM) facility. When the metadata is merged, the following inheritance rules are followed.

Table 20. Metadata merger inheritance rules

Scenario	Rules
Security metadata in deployment descriptor only	No merge is needed, the security metadata from the deployment descriptor is propagated.
Security metadata in annotations only	No merge is needed, the security metadata defined with annotations is propagated.
Security metadata in deployment descriptor and annotations	The metadata from the deployment descriptor and annotations is merged. The metadata in annotations is overridden by the same type of data from the deployment descriptor.

Five security annotations are currently supported. For each annotation, a MergeAction implementation is defined.

- **@DeclareRoles** (Servlet 2.5 and EJB 3)

The MergeAction implementation finds all the classes annotated with the DeclareRoles annotation. Within each annotated class for each role name specified, if the security roles listed in the deployment descriptor does not contain a SecurityRole with the annotated role name, a new SecurityRole is created and added to this list of security roles.

- **@RunAs** (Servlet 2.5 and EJB 3)

The MergeAction implementation finds all the classes with the RunAs annotation. For each annotated class, it finds the Servlet or the Enterprise Java Bean (EJB) associated with the given class. It then determines if a run-as element is defined in the deployment descriptor for the servlet or EJB. If one is not found, a new run-as element is created and added to the deployment descriptor. If a run-as element is found, this run-as element will be used instead of creating a new one. The role name used in the RunAs annotation must be defined in the deployment descriptor.

- **@DenyAll** (EJB 3 only)

The MergeAction implementation finds all the methods annotated with the DenyAll annotation. For each annotated method, if the method is not included in the deployment descriptor list of excluded methods, and a MethodPermission does not exist in the deployment descriptor, a new MethodElement is created and added to this list of excluded methods in the deployment descriptor.

- **@PermitAll** (EJB 3 only)

The MergeAction implementation finds all the classes and the methods with the PermitAll annotation. For each annotated class, it finds the Enterprise Java Bean (EJB) associated with the given class. It then searches the subset of the MethodElements in the list of all the MethodPermissions defined in the deployment descriptor for this EJB. If a MethodElement with a wildcard method name ("*") is not found and a wildcard method does not exist in the list of excluded methods or in the list of MethodElements with security roles, a new MethodPermission and a new MethodElement are created. The new

MethodPermission is marked unchecked and is added to the MethodPermission list in the deployment descriptor. The new MethodElement is added to the MethodElement list of the newly created unchecked MethodPermission. Similar action is done for all annotated methods. Instead of a wildcard MethodElement, the method signature must match exactly the signature of the annotated method.

- @RolesAllowed (EJB 3 only)

The MergeAction implementation finds all the classes and methods with the RolesAllowed annotation. For each annotated class, it finds the EJB associated with the given class. It then finds the subset of the MethodElements in the list of all the MethodPermissions defined in the deployment descriptor for this EJB. If a MethodElement with a wildcard method name (“*”) is not found, and a wildcard method does not exist in the list of excluded methods or in the list of unchecked MethodElements, a new MethodPermission and MethodElement are created. If a MethodPermission for this EJB exists with exactly the same roles as those found in the annotation, this MethodPermission will be used instead of creating a new one. For each role name specified in the annotation, a new SecurityRole is created and added to the SecurityRole list in the MethodPermission, If the MethodPermission was newly created, it is added to the MethodPermission list in the deployment descriptor. The new MethodElement created is added to the MethodElement list of the MethodPermission. Similar processing is done for all annotated methods. Instead of a wildcard MethodElement, the method signature must exactly match the signature of the annotated method. Additionally, for each role name specified in the annotation, if the deployment descriptor list of security roles does not contain a SecurityRole with the annotated role name, a new SecurityRole is also created and added to this list of security roles.


Delegations

Delegation is a process security identity propagation from a caller to a called object. As per the Java Platform, Enterprise Edition (Java EE) specification, a servlet and enterprise beans can propagate either the client or remote user identity when invoking enterprise beans, or they can use another specified identity as indicated in the corresponding deployment descriptor.

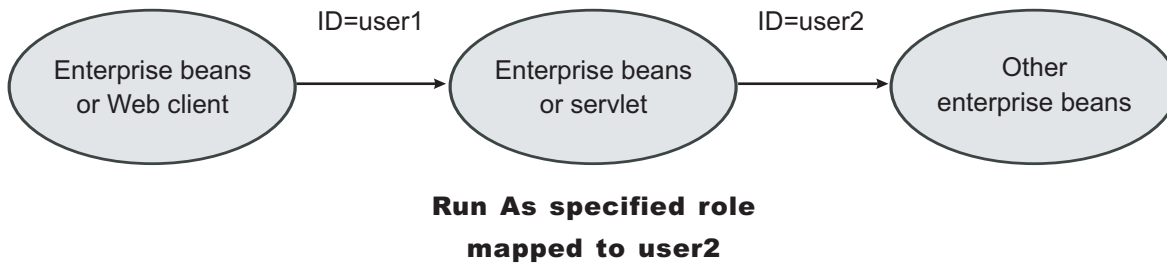
The extension supports enterprise bean propagation to the server ID when invoking other entity beans. Three types of delegations are possible:

- Delegate (RunAs) client identity
- Delegate (RunAs) specified identity
- Delegate (RunAs) system identity


Delegate (RunAs) client identity



Delegate (RunAs) specified identity



Delegate (RunAs) system identity



Note: The RunAs system identity delegation only works when server ID and password are used. When the `internalServerId` feature is used, it does not work because runAs with system identity is not supported. You must specify RunAs roles. When `internalServerId` is used, use the `RunAsSpecified` with a user ID and password that is mapped to the administrator role. See “Administrative roles and naming service authorization” on page 510 for more information about `internalServerId`.

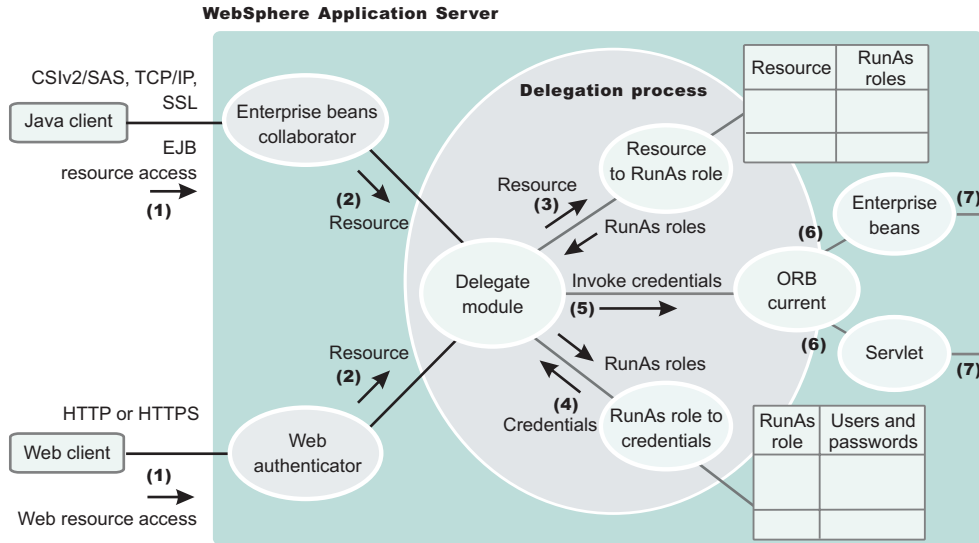
The EJB specification only supports delegation (RunAs) at the Enterprise JavaBeans (EJB) level. But an extension allows EJB method-level RunAs specification. With an EJB method level, the RunAs specification can specify a different RunAs role for different methods within the same enterprise beans.

The RunAs specification is detailed in the deployment descriptor, which is the `ejb-jar.xml` file in the EJB module and the `web.xml` file in the Web module. The extension to the RunAs specification is included in the `ibm-ejb-jar-ext.xml` file.

An IBM-specific binding file is available for each application that contains a mapping from the RunAs role to the user. This file is specified in the `ibm-application-bnd.xml` file.

These specifications are read by the runtime during application startup. The following figure illustrates the delegation mechanism, as implemented in the WebSphere Application Server security model.

Delegation



Delegation Process

Two tables help in the delegation process:

- Resource to RunAs role mapping table
- RunAs role to user ID and password mapping table

Use the Resource to RunAs role mapping table to get the role that is used by a servlet or by enterprise beans to propagate to the next enterprise beans call.

Use the RunAsRole to user ID and password mapping table to get the user ID that belongs to the RunAs role and its password.

Delegation is performed after successful authentication and authorization. During this process, the delegation module consults the Resource to RunAs role mapping table to get the RunAs role (3). The delegation module consults the RunAs role to user ID and password mapping table to get the user that belongs to the RunAs role (4). The user ID and password is used to create a new credential using the authentication module, which is not shown in the figure.

The resulting credential is stored in the Object Request Broker (ORB) Current as an invocation credential (5). Servlet and enterprise beans when invoking other enterprise beans pick up the invocation credential from the ORB Current (6) and call the next enterprise beans (7).

Authorizing access to J2EE resources using Tivoli Access Manager

The Java Authorization Contract for Containers (JACC) defines a contract between Java Platform, Enterprise Edition (Java EE) containers and authorization providers. You can use the default authorization or an external JACC authorization provider. When security is enabled in WebSphere Application Server, the default authorization is used unless a JACC provider is specified.

Before you begin

JACC enables any third-party authorization providers to plug into a Java EE application server (such as WebSphere Application Server) to make the authorization decisions when a Java EE resource is accessed. By default, WebSphere Application Server implements the JACC provider by using Tivoli Access Manager as the external authorization provider.

Read the following articles for more detailed information about JACC before you attempt to configure WebSphere Application Server to use a JACC provider:

- “JACC support in WebSphere Application Server” on page 523
- “JACC providers” on page 525
- “Tivoli Access Manager integration as the JACC provider” on page 529

Using the built-in authorization provider

You can extend the capabilities of WebSphere Application Server by plugging in your own authorization provider. You can use the built-in authorization or an external JACC authorization provider.

About this task

For an explanation of the administrative console panels that support these capabilities, see:

- Use the built-in authorization provider. It is recommended that you do not modify any settings on the authorization provider panels if you use the **Built-in authorization** option. For more information, see “External authorization provider settings.”
- Use an external authorization provider. If you use the **External authorization using a JACC provider** option, the external providers must be based on the Java Authorization Contract for Containers (JACC) specification to handle the Java Platform, Enterprise Edition (Java EE) authorization. By default, WebSphere Application Server enables you to configure the Tivoli Access Manager Java Authorization Contract for Containers (JACC) provider as the default external JACC provider. For more information, see “External Java Authorization Contract for Containers provider settings” on page 539 and “Tivoli Access Manager JACC provider settings” on page 546.

External authorization provider settings

Use this page to enable a Java Authorization Contract for Containers (JACC) provider for authorization decisions.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Click **External authorization providers**.

The application server provides a default authorization engine that performs all of the authorization decisions. In addition, the application server also supports an external authorization provider using the JACC specification to replace the default authorization engine for Java Platform, Enterprise Edition (Java EE) applications.

JACC is part of the Java EE specification, which enables third-party security providers such as Tivoli Access Manager to plug into the application server and make authorization decisions.

Note: Unless you have an external JACC provider or want to use a JACC provider for Tivoli Access Manager that can handle Java EE authorizations based on JACC, and it is configured and set up to use with the application server, do not enable **External authorization using a JACC provider**.

Built-in authorization:

Use this option all the time unless you want an external security provider such as the Tivoli Access Manager to perform the authorization decision for Java EE applications that are based on the JACC specification.

External JACC provider: Use this link to configure the application server to use an external JACC provider. For example, to configure an external JACC provider, the policy class name and the policy configuration factory class name are required by the JACC specification.

The default settings that are contained in this link are used by Tivoli Access Manager for authorization decisions. If you intend to use another provider, modify the settings as appropriate.

External Java Authorization Contract for Containers provider settings

Use this page to configure the application server to use an external Java Authorization Contract for Containers (JACC) provider. For example, the policy class name and the policy configuration factory class name are required by the JACC specification.

Use these settings when you have set up an external security provider that supports the JACC specification to work with the application server. The configuration process involves installing and configuring the provider server and configuring the client of the provider in the application server to communicate with the server. If the JACC provider is not enabled, these settings will be ignored.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Click **External authorization providers**.
3. Under Authorization provider, click **External JACC provider**.

Use the default settings when you use Tivoli Access Manager as the JACC provider. Install and configure the Tivoli Access Manager server prior to using it with the application server. Use the Tivoli Access Manager properties link under Additional properties, and configure the Tivoli Access Manager client in the application server to use the Tivoli Access Manager server. If you intend to use another provider, modify the settings as appropriate.

Name:

Specifies the name that is used to identify the external JACC provider.

This field is required.

Data type: String

Description:

Provides an optional description for the provider.

Data type: String

Policy class name:

Specifies a fully qualified class name that represents the `javax.security.jacc.policy.provider` property as per the JACC specification. The class represents the provider-specific implementation of the `java.security.Policy` abstract methods.

The class file must reside in the class path of each application server process. These processes include the application server, node agents and the deployment manager.

Do not add the Java archive (JAR) file, which contains the class file, to the `<WAS_HOME>/lib` directory in a product environment as service releases overwrite files in this directory.

This class is used during authorization decisions. The default class name is for Tivoli Access Manager implementation of the policy file.

This field is required. For information on enabling the JACC provider using this field, see the "Enabling the JACC provider for Tivoli Access Manager" article in the information center.

Data type: String
Default: com.tivoli.pd.as.jacc.TAMPolicy

Policy configuration factory class name:

Specifies a fully qualified class name that represents the javax.security.jacc.PolicyConfigurationFactory.provider property as per the JACC specification. The class represents the provider-specific implementation of the javax.security.jacc.PolicyConfigurationFactory abstract methods.

The class file must reside in the class path of each application server process. These processes include the application server, node agents and the deployment manager.

Do not add the Java archive (JAR) file, which contains the class file, to the <WAS_HOME>/lib directory in a product environment as service releases overwrite files in this directory.

This class represents the provider-specific implementation of the PolicyConfigurationFactory abstract class. This class is used to propagate the security policy information to the JACC provider during the installation of the J2EE application. The default class name is for the Tivoli Access Manager implementation of the policy configuration factory class name.

This field is required.

Data type: String
Default: com.tivoli.pd.as.jacc.TAMPolicyConfigurationFactory

Role configuration factory class name:

Specifies a fully qualified class name that implements the com.ibm.wsspi.security.authorization.RoleConfigurationFactory interface.

The class file must reside in the class path of each application server process. These processes include the application server, node agents and the deployment manager.

Do not add the Java archive (JAR) file, which contains the class file, to the <WAS_HOME>/lib directory in a product environment as service releases overwrite files in this directory.

When you implement this class, the authorization table information in the binding file is propagated to the provider during the installation of the J2EE application. The default class name is for the Tivoli Access Manager implementation of the role configuration factory class name.

This field is optional. For information on enabling the JACC provider using this field, see the "Enabling the JACC provider for Tivoli Access Manager" article in the information center.

Data type: String
Default: com.tivoli.pd.as.jacc.TAMRoleConfigurationFactory

Provider initialization class name:

Specifies a fully qualified class name that implements the com.ibm.wsspi.security.authorization.InitializeJACCProvider interface.

The class file must reside in the class path of each application server process. These processes include the application server, node agents and the deployment manager.

Do not add the Java archive (JAR) file, which contains the class file, to the <WAS_HOME>/lib directory in a product environment as service releases overwrite files in this directory.

When implemented, this class is called at the start and the stop of all the application server processes. You can use this class for any required initialization that is needed by the provider client code to communicate with the provider server. The properties that are entered in the custom properties link are passed to the provider when the process starts up. The default class name is for the Tivoli Access Manager implementation of the provider initialization class name.

This field is optional. For information on enabling the JACC provider using this field, see the "Enabling the JACC provider for Tivoli Access Manager" article in the information center.

Data type:	String
Default:	com.tivoli.pd.as.jacc.cfg.TAMConfigInitialize

Requires the EJB arguments policy context handler for access decisions:

Specifies whether the JACC provider requires the EJBArgumentsPolicyContextHandler handler to make access decisions.

Because this option has an impact on performance, do not set it unless it is required by the provider. Normally, this handler is required only when the provider supports instance-based authorization. Tivoli Access Manager does not support this option for J2EE applications.

Default:	Disabled
-----------------	----------

Supports dynamic module updates:

Specifies whether you can apply changes made to security policies of Web modules in a running application, dynamically without affecting the rest of the application.

If this option is enabled, the security policies of the added or modified Web modules are propagated to the JACC provider and only the affected Web modules are started.

If this option is disabled, then the security policies of the entire application are propagated to the JACC provider for any module-level changes. The entire application is restarted for the changes to take effect.

Typically, this option is enabled for an external JACC provider.

Default:	Enabled
-----------------	---------

Custom properties:

Specifies the properties that are required by the provider.

These properties are propagated to the provider during the startup process when the provider initialization class name is initialized. If the provider does not implement the provider initialization class name as described previously, the properties are not used.

The Tivoli Access Manager implementation does not require that you enter any properties in this link.

Tivoli Access Manager properties:

Specifies properties that are required by the Tivoli Access Manager implementation.

These properties are used to set up the communication between the application server and the Tivoli Access Manager server. You must install and configure the Tivoli Access Manager server before entering these properties.

Enabling an external JACC provider

Use this topic to enable an external JACC provider using the administrative console.

Before you begin

The Java Authorization Contract for Containers (JACC) defines a contract between Java Platform, Enterprise Edition (Java EE) containers and authorization providers. This contract enables any third-party authorization providers to plug into a Java EE 5 application server, such as WebSphere Application Server to make the authorization decisions when a Java EE resource is accessed.

1. From the WebSphere Application Server administrative console, click **Security > Global security > External authorization providers**.
2. Under Related items, click **External JACC provider**.
3. The fields are set for Tivoli Access Manager by default. If you do not plan to use Tivoli Access Manager as the JACC provider, replace these fields with the details for your own external JACC provider.
4. If any custom properties are required by the JACC provider, click **Custom properties** under Additional properties and enter the properties. When using the Tivoli Access Manager, use the **Tivoli Access Manager properties** link instead of the Custom properties link. For more information, see “Configuring the JACC provider for Tivoli Access Manager using the administrative console.”
5. On the External authorization providers panel, select the **External authorization using a JACC provider** option and click **OK**.
6. Complete the remaining steps to enable security. If you are using Tivoli Access Manager, you must select LDAP as the user registry and use the same LDAP server. For more information on configuring LDAP registries, see “Configuring Lightweight Directory Access Protocol user registries” on page 119.
7. Verify that all of the changes are synchronized across all nodes. For more information, see Synchronizing nodes with the wsadmin tool
8. In a multinode environment, stop and start the deployment manager configuration.

Issue the following commands:

```
profile_root/bin/stopManager.bat  
-username user_name  
-password password
```

```
profile_root/bin/startManager.bat
```

9. Restart all servers to make these changes effective.

Configuring the JACC provider for Tivoli Access Manager using the administrative console

Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

Before you begin

Before configuring Tivoli Access Manager as the JACC provider, verify that all of the managed servers, including node agents, are started

Prior to completing the following steps, verify that you have previously created a security administrative user. For more information, see “Creating the security administrative user” on page 544.

About this task

The following configuration is performed on the management server. When you click either **Apply** or **OK**, configuration information is checked for consistency, saved, and applied if successful.

This configuration information is propagated to the nodes when synchronization is performed. Restart the nodes for the configuration changes to take effect.

To configure Tivoli Access Manager as the JACC provider using the administrative console, complete the following steps:

1. Start the WebSphere Application Server administrative console by clicking `http://yourhost.domain:port_number/ibm/console` after starting WebSphere Application Server. If security is currently disabled, log in with any user ID. If security is currently enabled, log in with a predefined administrative ID and password. This ID is typically the server user ID that is specified when you configure the user registry.
2. Click **Security > Global security > External authorization providers**.
3. Under General properties, select **External authorization using a JACC provider**.
4. Under Related items, click **External JACC provider**.
5. Under Additional properties, click **Tivoli Access Manager Properties**. The Tivoli Access Manager JACC provider configuration screen is displayed.
6. Enter the following information:

Enable embedded Tivoli Access Manager

Select this option to enable Tivoli Access Manager.

Ignore errors during embedded Tivoli Access Manager disablement

Select this option when you want to unconfigure the JACC provider. Do not select this option during configuration.

Client listening port set

WebSphere Application Server must listen using a TCP/IP port for authorization database updates from the policy server. More than one process can run on a particular node or machine. More than one authorization server can be specified by separating the entries with commas. Specifying more than one authorization server at a time is useful for reasons of failover and performance. Enter the listening ports used by Tivoli Access Manager clients, separated by a comma. If a range of ports is specified, separate the lower and higher values by a colon (:) (for example, 7999, 9990:999).

Policy server

Enter the name of the Tivoli Access Manager policy server and the connection port. Use the `policy_server:port` form. The policy communication port is set at the time of the Tivoli Access Manager configuration, and the default is 7135.

Authorization servers

Enter the name of the Tivoli Access Manager authorization server. Use the `auth_server:port:priority` form. The authorization server communication port is set at the time of the Tivoli Access Manager configuration, and the default is 7136. The priority value is determined by the order of the authorization server use (for example, `auth_server1:7136:1` and `auth_server2:7137:2`). A priority value of 1 is required when configuring against a single authorization server.

Administrator user name

Enter the Tivoli Access Manager administrator user name that was created when Tivoli Access Manager was configured; it is usually `sec_master`.

Administrator user password

Enter the Tivoli Access Manager administrator password.

User registry distinguished name suffix

Enter the distinguished name suffix for the user registry that is shared between Tivoli Access Manager and WebSphere Application Server, for example, `o=ibm, c=us`.

Security domain

You can create more than one security domain in Tivoli Access Manager, each with its own administrative user. Users, groups and other objects are created within a specific domain, and are not permitted to access resource in another domain. Enter the name of the Tivoli Access Manager security domain that is used to store WebSphere Application Server users and groups.

If a security domain is not established at the time of the Tivoli Access Manager configuration, leave the value as `Default`.

Administrator user distinguished name

Enter the full distinguished name of the WebSphere Application Server security administrator ID (for example, `cn=wasadmin, o=organization, c=country`). The ID name must match the Server user ID on the Lightweight Directory Access Protocol (LDAP) User Registry panel in the administrative console. To access the LDAP User Registry panel, click **Security > Global security**. Under **User account repository**, choose **Standalone LDAP registry** as the available realm definition. Then click **Configure**.

7. When all information is entered, click **OK** to save the configuration properties. The configuration parameters are checked for validity and the configuration is attempted at the host server or cell manager.

Results

After you click **OK**, WebSphere Application Server completes the following actions:

- Validates the configuration parameters.
- Configures the host server or cell manager.

These processes might take some time depending on network traffic or the speed of your machine.

What to do next

If the configuration is successful, the parameters are copied to all subordinate servers, including the node agents. To complete the embedded Tivoli Access Manager client configuration, you must restart all of the servers, including the host server, and enable WebSphere Application Server security.

Creating the security administrative user:

Enabling security requires the creation of a WebSphere Application Server administrative user. Use the Tivoli Access Manager command-line `pdadmin` utility to create the Tivoli Access Manager administrative user for WebSphere Application Server. This utility is available on the policy server host machine.

About this task

Follow these steps to use the `pdadmin` utility.

1. From a command line, start the `pdadmin` utility as the Tivoli Access Manager administrative user, `sec_master`:

```
pdadmin -a sec_master -p sec_master_password
```

2. Create a WebSphere Application Server security user. For example, the following instructions create a new user, `wasadmin`. The command is entered as one continuous line:

```
pdadmin> user create wasadmin cn=wasadmin,o=organization,  
c=country wasadmin wasadmin myPassword
```

Substitute values for organization and country that are valid for your Lightweight Directory Access Protocol (LDAP) user registry.

3. Enable the account for the WebSphere Application Server security administrative user by issuing the following command:

```
pdadmin> user modify wasadmin account-valid yes
```

What to do next

Configure the Java Authorization Contract for Container (JACC) provider for Tivoli Access Manager. For more information, see “Tivoli Access Manager JACC provider configuration.”

Tivoli Access Manager JACC provider configuration:

You can configure the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager to deliver authentication and authorization protection for your applications or for authentication only. Most deployments that use the JACC provider for Tivoli Access Manager to configure Tivoli Access Manager provide both authentication and authorization functionality.

If you want Tivoli Access Manager to provide authentication, but leave authorization as part of WebSphere Application Server’s native security, add the `com.tivoli.pd.as.amwas.DisableAddAuthorizationTableEntry=true` property to the `amwas.amjacc.template.properties` file. The file is located in the `profile_root/config/cells/cell_name` directory.

After this property is set, perform the tasks for setting Tivoli Access Manager Security, as documented.

You can configure the JACC provider for Tivoli Access Manager using either the WebSphere Application Server administrative console or the **wsadmin** command-line utility.

- For details on configuring the JACC provider for Tivoli Access Manager using the administrative console, refer to “Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542.
- For details on configuring the Tivoli Access Manager JACC provider using the **wsadmin** command line utility, refer to “Configuring the JACC provider for Tivoli Access Manager using the wsadmin utility” on page 955.

The JACC configuration files for Tivoli Access Manager that are common across multiple WebSphere Application Server profiles are created by default under the `java/jre` directory. When you install WebSphere Application Server, you are given permissions to read and write to the files in this directory.

UNIX Profiles created by users who are different to the user that installed the application have read-only permissions for this directory.

This situation is not ideal because configuration of the JACC provider for Tivoli Access Manager fails in these situations. To avoid this situation, you can add the following property to the `profile_root/config/cells/cell_name/amwas.amjacc.template.properties` file: `com.tivoli.pd.as.jacc.CommonFileLocation=new location` where `new location` is a fully qualified directory name.

This property applies read and write permissions to the `java/jre` directory.

The **wsadmin** command is available to reconfigure the Java Authorization Contract for Containers (JACC) Tivoli Access Manager interface:

```
$AdminTask reconfigureTAM -interactive
```

This command effectively prompts you through the process of unconfiguring the interface and then reconfiguring it.

Tivoli Access Manager JACC provider settings:

Use this page to configure the Java Authorization Contract for Container (JACC) provider for Tivoli Access Manager.

Note: When a third-party authorization such as Tivoli Access Manager or SAF for z/OS is used, the information in the administrative console panel might not represent the data in the provider. Also, any changes to the panel might not be reflected in the provider automatically. Follow the provider's instructions to propagate any changes made to the provider.

To view the JACC provider settings for Tivoli Access Manager, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, click **External authorization providers**.
3. Under Authorization provider, click **External JACC provider**.
4. Click **Configure** to configure the properties for Tivoli Access Manager.

Enable embedded Tivoli Access Manager:

Enables or disables the embedded Tivoli Access Manager client configuration.

Default: Disabled
Range: Enabled or Disabled

Note: If you want to disable Tivoli Access Manager as the JACC provider, clear this option and also select **Default authorization**.

Ignore errors during embedded Tivoli Access Manager disablement:

When selected, errors are ignored during disablement of the embedded Tivoli Access Manager client.

This option is applicable only when re-configuring an embedded Tivoli Access Manager client or disabling an embedded Tivoli Access Manager.

Default: Disabled
Range: Enabled or Disabled

Client listening port set:

Enter the ports that are used as listening ports by Tivoli Access Manager clients.

The application server needs to listen on a TCP/IP port for authorization database updates from the policy server. More than one process can run on a particular node and machine, so a list of ports is required for use by the processes. If you specify a range of ports, separate the lower and higher values by a colon (:). Single ports and port ranges are specified on separate lines. An example list might look like the following example:

```
7999
8900:8999
```

Note: Each of the servants might need to open up a listener port.

Policy server:

Enter the name, fully-qualified domain name, or IP address of the Tivoli Access Manager policy server and the connection port.

Use the form *policy_server:port*. The policy server communication port was set at the time of the Tivoli Access Manager configuration. The default is 7135.

Authorization servers:

Enter the name, fully-qualified domain name, or IP address of the Tivoli Access Manager authorization server. Use the form, *auth_server:port:priority*.

The authorization server communication port is set at the time of Tivoli Access Manager configuration. The default is 7136. You can specify more than one authorization server by entering each server on a new line. Configuring more than one authorization server provides for failover. The priority value is the order of authorization server use. For example:

```
auth_server1.mycompany.com:7136:1  
auth_server2.mycompany.com:7137:2
```

A priority of 1 is still required when configuring a single authorization server.

Administrator user name:

Enter the Tivoli Access Manager administration user ID, as created at the time of Tivoli Access Manager configuration. This ID is usually, *sec_master*.

Administrator user password:

Enter the Tivoli Access Manager administration password for the user ID that is entered in the **Administrator user name** field.

User registry distinguished name suffix:

Enter the distinguished name suffix for the user registry to share between Tivoli Access Manager and the application server. For example: *o=organization,c=country*

Security domain:

Enter the name of the Tivoli Access Manager security domain that is used to store application server users and groups.

Specification of the Tivoli Access Manager domain is required because more than one security domain can be created in Tivoli Access Manager with its own administrative user. Users, groups, and other objects are created within a specific domain and are not permitted to access resources in another domain. If a security domain is not established at the time of Tivoli Access Manager configuration, leave the value as *Default*.

Default: Default

Administrator user distinguished name:

Enter the fully distinguished name of the security administrator ID for the application server. For example, *cn=wasadmin,o=organization,c=country*

JACC provider configuration properties for Tivoli Access Manager:

The JACC provider configuration properties detailed below may require configuration.

The Java property files are created in the *profile_root/etc/tam* directory.

Two properties files might require configuration:

- *amwas.node_name_server_name.amjacc.properties* contains properties that are used by the JACC provider of Tivoli Access Manager.
- *amwas.node_name_server_name.pdjlog.properties* contains logging properties that are created from the *amwas.pdjlog.template.properties* file for the specific node and server combination at the time of configuration.

Use *amwas.node_name_server_name.amjacc.properties* file to configure static role caching, dynamic role caching, object caching, and role-based policy framework properties.

Static role caching properties:

The static role cache holds role memberships that do not expire.

These properties are in the *profile_root/etc/tam/amwas.node_name_server_name.amjacc.properties* file.

The *profile_root* directory is the value of the *profilePath* parameter at profile creation time.

Related tasks

“Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542
Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

com.tivoli.pd.as.cache.EnableStaticRoleCaching=true:

Enables or disables static role caching. Static role caching is enabled by default.

com.tivoli.pd.as.cache.StaticRoleCache=com.tivoli.pd.as.cache.StaticRoleCacheImpl:

This property holds the implementation class of the static role cache. You do not need to change this property, although the opportunity exists to implement your own cache, if necessary.

com.tivoli.pd.as.cache.StaticRoleCache.Roles=Administrator,Operator,Monitor,Deployer:

Defines the administration roles for WebSphere Application Server.

Note: Enhance Application performance by adding the static roles: **CosNamingRead**, **CosNamingWrite**, **CosNamingCreate**, **CosNamingDelete**. These roles support for improved lookup performance within the application naming service.

Dynamic role caching properties:

The dynamic role cache holds role memberships that expire.

These properties are in the *profile_root/etc/tam/amwas.node_name_server_name.amjacc.properties* file.

The *profile_root* directory is the value of the *profilePath* parameter at profile creation time.

Related tasks

“Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542
Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

com.tivoli.pd.as.cache.EnableDynamicRoleCaching=true:

Enables or disables dynamic role caching. Dynamic role caching is enabled by default.

com.tivoli.pd.as.cache.DynamicRoleCache=com.tivoli.pd.as.cache.DynamicRoleCacheImpl:

This property holds the implementation class of the dynamic role cache. You do not need to change this property, although the opportunity exists to implement your own cache, if necessary.

com.tivoli.pd.as.cache.DynamicRoleCache.MaxUsers=100000:

The maximum number of users that the cache supports before a cache cleanup is performed. The default number of users is 100000.

com.tivoli.pd.as.cache.DynamicRoleCache.NumBuckets=20:

The number of tables that is used internally by the dynamic role cache. The default is 20. When a large number of threads use the cache, increase the value to tune and optimize cache performance.

com.tivoli.pd.as.cache.DynamicRoleCache.PrincipalLifeTime=10:

The period of time in minutes that a principal entry is stored in the cache. The default time is 10 minutes. The term, *principal*, here refers to the Tivoli Access Manager credential that is returned from a unique Lightweight Directory Access Protocol user.

com.tivoli.pd.as.cache.DynamicRoleCache.RoleLifetime=20:

The period of time in seconds that a role is stored in the role list for a user before it is discarded. The default is 20 seconds.

Object caching properties:

The object cache is used to cache all Tivoli Access Manager objects, including their extended attributes. This bypasses the need to query the Tivoli Access Manager authorization server for each resource request.

These properties are in the *profile_root/etc/tam/amwas.node_name_server_name.amjacc.properties* file.

The *profile_root* directory is the value of the *profilePath* parameter when the profile is created.

These object cache properties cannot be changed after configuration. If any require changing, it should be done before configuration of the nodes in the cell. Changes need to be made in the template properties file before any configuration actions are performed. Properties changed after configuration might cause access decisions to fail.

Related tasks

“Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542
Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

com.tivoli.pd.as.cache.EnableObjectCaching=true:

This property enables or disables object caching. The default value is true.

com.tivoli.pd.as.cache.ObjectCache=com.tivoli.pd.as.cache.ObjectCacheImpl:

This property is the class used to perform object caching. You can implement your own object cache if required. This can be done by implementing the *com.tivoli.pd.as.cache.IObjectCache* interface. The default is *com.tivoli.pd.as.cache.ObjectCacheImpl*.

com.tivoli.pd.as.cache.ObjectCache.NumBuckets=20:

This property specifies the number of buckets used to store object cache entries in the underlying hash table. The default is 20.

com.tivoli.pd.as.cache.ObjectCache.MaxResources=10000:

This property specifies the total number of entries for all buckets in the cache. This figure, divided by NumBuckets determines the maximum size of each bucket. The default is 10000.

com.tivoli.pd.as.cache.ObjectCache.ResourceLifeTime=20:

This property specifies the length of time in minutes that objects are kept in the object cache. The default is 20.

Role-based policy framework properties:

Although it is very unlikely that you will need to change these properties, use this file to reference supported properties within the role-based policy framework.

The role-based policy framework parameters are located in the Java Authorization Contract for Containers (JACC) configuration file and in the authorization configuration file. They are set at the time of JACC provider configuration and authorization server configuration. The role-based policy framework settings for the authorization table and the JACC provider can be modified separately for each WebSphere Application Server instance. The *amwas.node_server.authztable.properties* configuration file is generated from the authorization table. The *amwas.node_name_server_name.amjacc.properties* configuration file is generated from the JACC provider. Both files are stored in the *profile_root/etc/tam* directory. It is very unlikely that you might need to change these properties. The properties are described here for reference.

The settings cannot be changed after configuration. Make changes in the template properties file before any configuration actions are performed. Properties that are changed after configuration will cause access decisions to fail.

Related tasks

“Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542
Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

com.tivoli.pd.as.rbpf.AMAAction=i:

This property is used to signify that a user is granted access to a role. This value is added to a Tivoli Access Manager access control list (ACL) and places invoke access on roles for users and groups.

com.tivoli.pd.as.rbpf.AMActionGroup=WebAppServer:

This property sets the Tivoli Access Manager action group that serves as a container for the action that is specified by the `com.tivoli.pd.as.rbpf.AMAction` property. The permission set in the `com.tivoli.pd.as.rbpf.AMAction` property goes into this action group.

com.tivoli.pd.as.rbpf.PosRoot=WebAppServer:

This property is used to determine where roles are stored in the protected object space.

com.tivoli.pd.as.rbpf.ProductId=deployedResources:

This property specifies the location under the root location that is specified in the `posroot` property to separate other products in the protected object space. Embedded Tivoli Access Manager objects are found in the `/WebAppServer/deployedResources` directory. The default value is `deployedResources`.

com.tivoli.pd.as.rbpf.ResourceContainerName=Resources:

This property specifies the Tivoli Access Manager object space container name for the protected resources. The default location is the `/WebAppServer/deployedResources/Resources` directory.

com.tivoli.pd.as.rbpf.RoleContainerName=Roles:

This property specifies the Tivoli Access Manager protected object space container name for the security roles. The default location is the `/WebAppServer/deployedResources/Roles` directory.

System-dependent configuration properties:

Do not change these system-dependent configuration properties. These properties are included in this article for reference only.

These properties are in the `app_server_root/etc/amwas.node_name_server_name.amjacc.properties` file.

The `profile_root` variable is the value of the `profilePath` parameter when the profile is created.

Related tasks

“Configuring the JACC provider for Tivoli Access Manager using the administrative console” on page 542
Use this task to configure Tivoli Access Manager as the Java Authorization Contract for Containers (JACC) provider using the administrative console.

com.tivoli.pd.as.rbpf.AmasSession.CfgURL=file/:\$WAS_HOME/profiles/profile_name/etc/tam/amwas.node_server.pdperm.properties:

This entry is generated by the Java Authorization Contract for Containers (JACC) provider configuration. This argument specifies the location of the file that contains information about the JACC provider of Tivoli Access Manager. Do not change this entry or the properties in the `amwas.node_server.pdperm.properties` file.

com.tivoli.pd.as.rbpf.AmasSession.CfgURL=file:/user_root/etc/tam/amwas.node_server.pdperm.properties:

This entry is generated by the Java Authorization Contract for Containers (JACC) provider configuration. It specifies the location of the file that contains information about the Tivoli Access Manager JACC provider. Do not change this entry or the properties in the `amwas.node_server.pdperm.properties` file.

com.tivoli.pd.as.rbpf.AmasSession.LoggingURL=file/:\$WAS_HOME/profiles/profile_name/etc/tam/amwas.node_server.pdlog.properties:

This entry contains the location of the logging configuration file for the JACC provider of Tivoli Access Manager. The referenced file is generated by the JACC provider of Tivoli Access Manager configuration. Do not change this entry.

*com.tivoli.pd.as.rbpf.AmasSession.LoggingURL=file:/user_root/etc/tam/
amwas.node_server.pdjlog.properties:*

This entry contains the location of the logging configuration file for the Tivoli Access Manager JACC provider. The file referenced is generated by the Tivoli Access Manager JACC provider configuration. Do not change this entry.

Administering security users and roles with Tivoli Access Manager

Use these steps to manage user-to-role mappings and user-to-group mappings for applications.

About this task

User-to-role mapping and user-to-group mapping for the JACC provider of Tivoli Access Manager are performed using the WebSphere Application Server administrative console.

1. Click **Applications > Enterprise applications > application_name**.
2. Under Additional properties, click **Security role to user/group mapping**. The user and groups management screen is displayed.
3. Select the role that requires user or group management and use **Lookup users** or **Lookup groups** to manage the users or groups for the selected role. The native role mapping uses the MapRolesToUsers administrative task. If you are using Tivoli Access Manager, use the TAMMapRolesToUsers administrative task instead. The syntax and options for the Tivoli version are the same as those used in the native version. For more information, see `csec_role_based_sec.dita` and `tsec_use_TAM_groups.dita`.

Configuring Tivoli Access Manager groups

Use these steps to configure the WebSphere Application Server administrative console to add objects of the `accessGroup` class to the list of object classes that represent user registry groups.

About this task

You can use the WebSphere Application Server administrative console to specify security policies for applications that run in the WebSphere Application Server environment. You can also use the WebSphere Application Server administrative console to specify security policies for other Web resources, based on the entities that are stored in the user registry.

Tivoli Access Manager adds the `accessGroup` object class to the registry. Tivoli Access Manager administrators can use the `pdadmin` utility, which is available only on the policy server host in the `PD.RTE` fileset, to create new groups. These new groups are added to the registry as the `accessGroup` object class.

The WebSphere Application Server administrative console is not configured by default to recognize objects of the `accessGroup` class as user registry groups. You can configure the WebSphere Application Server administrative console to add this object class to the list of object classes that represent user registry groups. To do this configuration, complete the following instructions:

1. From the WebSphere Application Server administrative console, access the advanced settings for configuring security by clicking **Security > Global security**.
2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry**, and click **Configure**.
3. Under Additional properties, click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.

4. Modify the **Group Filter** field. Add the following entry: (objectclass=accessGroup)
The Group Filter field looks like the following example:

```
(&(cn=%w)(|(objectclass=groupOfNames)
(objectclass=groupOfUniqueNames)(objectclass=accessGroup)))
```

5. Modify the **Group Member ID Map** field. Add the following entry: accessGroup:member. The Group Member ID Map field looks like the following example:

```
groupOfNames:member;groupOfUniqueNames:uniqueMember;
accessGroup:member
```

6. Stop and restart WebSphere Application Server.

Configuring additional authorization servers

Tivoli Access Manager secure domains can contain more than one authorization server. Having multiple authorization servers is useful for providing a failover capability as well as improving performance when the volume of access requests is large.

1. Refer to the *Tivoli Access Manager Base Administration Guide* for details on installing and configuring authorization servers. This document is available in the IBM Tivoli Access Manager for e-business information center.
2. Re-configure the Java Authorization Contract for Containers (JACC) provider using the \$AdminTask reconfigureTAM interactive **wsadmin** command. Enter all new and existing options.

Logging Tivoli Access Manager security

Use this topic to enable the trace specification to indicate tracing at the required level.

About this task

The Java Authorization Contract for Containers (JACC) for Tivoli Access Manager provider messages are logged to the configured trace output location, and messages are written to standard out SystemOut.log file. When trace is enabled, all logging, both trace and messaging, is sent to the trace.log file.

1. The `amwas.node_server.pdjournal.properties` file must be updated and the **isLogging** attribute set to *true* for the required component. For example, to enable tracing for the JACC provider for Tivoli Access Manager, set the following line to *true*:
`amwas.node_server.pdjournal.properties:baseGroup.AMWASWebTraceLogger.isLogging=true`
2. Enable tracing for the JACC provider of Tivoli Access Manager components in the WebSphere Application Server administrative console by completing the following steps:
 - a. Click **Troubleshooting > Logs and Trace > server_name**.
 - b. Under Logs and Trace tasks, click **Diagnostic trace**.
 - c. Select the **Enable Log** option.
 - d. Click **Apply**.
 - e. Click **Troubleshooting > Logs and Trace > server_name**.
 - f. Click **Change Log Detail Levels**.
 - g. Click **Components**. Tracing for all components can be enabled using the **com.tivoli.pd.as.*** command. Tracing for separate components can be enabled using the following commands:
 - **com.tivoli.pd.as.rbpf.*** for role-based policy framework tracing
 - **com.tivoli.pd.as.jacc.*** for JACC provider tracing
 - **com.tivoli.pd.as.pdwas.*** for the authorization table
 - **com.tivoli.pd.as.cfg.*** for configuration
 - **com.tivoli.pd.as.cache.*** for caching

For more information, see `utrb_loglevel.dita`.

- h. Click **Apply**.

What to do next

The trace specification now indicates that tracing is enabled at the required level. Save the configuration and restart the server for the changes to take effect.

Related reference

Diagnostic trace service settings

Tivoli Access Manager loggers:

The Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager uses the JLog logging framework as does the Java runtime environment for Tivoli Access Manager. You can enable tracing and messaging selectively for specific JACC providers for Tivoli Access Manager components.

The JACC for Tivoli Access Manager provider messages are logged to the configured trace output location, and messages are written to standard out `SystemOut.log` file. When trace is enabled, all logging, both trace and messaging, is sent to the `trace.log` file.

Tracing and message logging for the JACC provider for Tivoli Access Manager are configured in the `amwas.node_server.pdolog.properties` properties file, which is located in the `profile_root/etc/tam` directory. This file contains logging properties from the `amwas.pdolog.template.properties` template file for the specific node and server combination at the time of JACC provider for Tivoli Access Manager configuration.

The contents of this file let the user control:

- Whether tracing is enabled or disabled for the JACC provider of Tivoli Access Manager components.
- Whether message logging is enabled or disabled for the JACC provider of Tivoli Access Manager components.

The `amwas.node_server.pdolog.properties` file defines several loggers, each of which is associated with one JACC provider of Tivoli Access Manager component. These loggers include:

Logger Name	Description
AmasRBPFTTraceLogger AmasRBPFMMessageLogger	Logs messages and trace for the role-based policy framework. This underlying framework is used by embedded Tivoli Access Manager to make access decisions.
AmasCacheTraceLogger AmasCacheMessageLogger	Logs messages and trace for the policy caches that are used by the role-based policy framework.
AMWASWebTraceLogger AMWASWebMessageLogger	Logs messages and trace for the WebSphere Application Server authorization plug-in.
AMWASConfigTraceLogger AMWASConfigMessageLogger	Logs messages and trace for the configuration actions of the JACC provider for Tivoli Access Manager .
JACCTraceLogger JACCMessageLogger	Logs messages and trace for the JACC provider activity of Tivoli Access Manager .

Note: Tracing can have a significant impact on system performance. Enable tracing only when diagnosing the cause of a problem.

The implementation of these loggers routes messages to the WebSphere Application Server logging sub-system. All messages are written to the WebSphere Application Server `trace.log` file.

For each logger, the `amwas.node_server.pdolog.properties` file defines an `isLogging` attribute which, when set to `true`, enables logging for the specific component. A value of `false` disables logging for that component.

The `amwas.node_server.pdolog.properties` file defines the parent loggers `MessageLogger` and `TraceLogger` that also have an `isLogging` attribute. If the child loggers do not specify this `isLogging` attribute, they inherit the value of their respective parent. When the JACC provider for Tivoli Access Manager is enabled, the `isLogging` attribute is set to `true` for the `MessageLogger` and set to `false` for the `TraceLogger` logger. Message logging is enabled for all components and tracing is disabled for all components, by default.

To turn on tracing for a JACC provider component, see [Logging Tivoli Access Manager security](#).

Interfaces that support JACC

WebSphere Application Server provides the `RoleConfigurationFactory` and the `RoleConfiguration` interfaces, which are similar to `PolicyConfigurationFactory` and `PolicyConfiguration` interfaces so the information that is stored in the bindings file can be propagated to the provider during installation. The implementation of these interfaces is optional.

RoleConfiguration interface

Use the `RoleConfiguration` interface to propagate the authorization information to the provider. This interface is similar to the `PolicyConfiguration` interface that is found in Java Authorization Contact for Containers (JACC).

```
RoleConfiguration
    - com.ibm.wsspi.security.authorization.RoleConfiguration

/**
 * This interface is used to propagate the authorization table information
 * in the binding file during application installation. Implementation of this interface is
 * optional. When a JACC provider implements this interface during an application, both
 * the policy and the authorization table information are propagated to the provider.
 * If this is not implemented, only the policy information is propagated as per
 * the JACC specification.
 * @ibm-spi
 * @ibm-support-class-A1
 */

public interface RoleConfiguration
/**
 * Add the users to the role in RoleConfiguration.
 * The role is created, if it does not exist in RoleConfiguration.
 * @param role the role name.
 * @param users the list of the user names.
 * @exception RoleConfigurationException if the users cannot be added.
 */
public void addUsersToRole(String role, List users)
    throws RoleConfigurationException

/**
 * Remove the users to the role in RoleConfiguration.
 * @param role the role name.
 * @param users the list of the user names.
 * @exception RoleConfigurationException if the users cannot be removed.
 */
public void removeUsersFromRole(String role, List users)
    throws RoleConfigurationException

/**
 * Add the groups to the role in RoleConfiguration.
 * The role is created if it does not exist in RoleConfiguration.
 * @param role the role name.
 * @param groups the list of the group names.
 * @exception RoleConfigurationException if the groups cannot be added.
 */
public void addGroupsToRole(String role, List groups)
    throws RoleConfigurationException
```

```

/**
 * Remove the groups to the role in RoleConfiguration.
 * @param role the role name.
 * @param groups the list of the group names.
 * @exception RoleConfigurationException if the groups cannot be removed.
 */
public void removeGroupsFromRole( String role, List groups)
throws RoleConfigurationException

/**
 * Add the everyone to the role in RoleConfiguration.
 * The role is created if it does not exist in RoleConfiguration.
 * @param role the role name.
 * @exception RoleConfigurationException if the everyone cannot be added.
 */
public void addEveryoneToRole(String role)
throws RoleConfigurationException

/**
 * Remove the everyone to the role in RoleConfiguration.
 * @param role the role name.
 * @exception RoleConfigurationException if the everyone cannot be removed.
 */
public void removeEveryoneFromRole( String role)
throws RoleConfigurationException

/**
 * Add the all authenticated users to the role in RoleConfiguration.
 * The role is created if it does not exist in RoleConfiguration.
 * @param role the role name.
 * @exception RoleConfigurationException if the authentication users cannot
 * be added.
 */
public void addAuthenticatedUsersToRole(String role)
throws RoleConfigurationException

/**
 * Remove the all authenticated users to the role in RoleConfiguration.
 * @param role the role name.
 * @exception RoleConfigurationException if the authentication users cannot
 * be removed.
 */
public void removeAuthenticatedUsersFromRole( String role)
throws RoleConfigurationException

/**
 * This commits the changes in Roleconfiguration.
 * @exception RoleConfigurationException if the changes cannot be
 * committed.
 */
public void commit( )
throws RoleConfigurationException

/**
 * This deletes the RoleConfiguration from the RoleConfiguration Factory.
 * @exception RoleConfigurationException if the RoleConfiguration cannot
 * be deleted.
 */
public void delete( )
throws RoleConfigurationException

/**
 * This returns the contextID of the RoleConfiguration.
 * @exception RoleConfigurationException if the contextID cannot be
 * obtained.
 */
public String getContextID( )
throws RoleConfigurationException

```

RoleConfigurationFactory interface

The RoleConfigurationFactory interface is similar to the PolicyConfigurationFactory interface that is introduced by JACC, and is used to obtain RoleConfiguration objects based on the contextID IDs.

```
RoleConfigurationFactory
- com.ibm.wsspi.security.authorization.RoleConfigurationFactory

/**
 * This interface is used to instantiate the com.ibm.wsspi.security.authorization.RoleConfiguration
 * objects based on the context identifier similar to the policy context identifier.
 * Implementation of this interface is required only if the RoleConfiguration interface is implemented.
 *
 * @ibm-spi
 * @ibm-support-class-A1
 */

public interface RoleConfigurationFactory
/**
 * This gets a RoleConfiguration with contextID from the
 * RoleConfigurationfactory. If the RoleConfiguration does not exist
 * for the contextID in the RoleConfigurationFactory, a new
 * RoleConfiguration with contextID is created in the
 * RoleConfigurationFactory. The contextID is similar to
 * PolicyContextID, but it does not contain the module name.
 * If remove is true, the old RoleConfiguration is removed and a new
 * RoleConfiguration is created, and returns with the contextID.
 * @return the RoleConfiguration object for this contextID
 * @param contextID the context ID of RoleConfiguration
 * @param remove true or false
 * @exception RoleConfigurationException if RoleConfiguration
 * cannot be obtained.
 */
public abstract com.ibm.ws.security.policy.RoleConfiguration
    getRoleConfiguration(String contextID, boolean remove)
        throws RoleConfigurationException
```

InitializeJACCProvider provider

When implemented by the provider, this interface is called by every process where the JACC provider can be used for authorization. All additional properties that are entered during the authorization check are passed to the provider. For example, the provider can use this information to initialize client code to communicate with their server or repository. The cleanup method is called during server shutdown to clean up the configuration.

Declaration

```
public interface InitializeJACCProvider
```

Description

This interface has two methods. The JACC provider can implement the interface, and WebSphere Application Server calls it to initialize the JACC provider. The name of the implementation class is obtained from the value of the initializeJACCProviderClassName system property.

This class must reside in a Java archive (JAR) file on the class path of each server that uses this provider.

```
InitializeJACCProvider
- com.ibm.wsspi.security.authorization.InitializeJACCProvider

/**
 * Initializes the JACC provider
 *   * @return 0 for success.
 * @param props the custom properties that are included for this provider will
 * pass to the implementation class.
 * @exception Exception for any problems encountered.
```



```

/**
public int initialize(java.util.Properties props)
throws Exception

/**
 * This method is for the JACC provider cleanup and will be called during a process stop.
 **/
public void cleanup()

```

Enabling the JACC provider for Tivoli Access Manager

The Java Authorization Contract for Container (JACC) provider for Tivoli Access Manager is configured by default. Use this topic to enable the JACC provider for Tivoli Access Manager.

About this task

Note: Do not perform this task if you are configuring the JACC provider for Tivoli Access Manager to supply authentication services only. Only perform this task for installations that require both Tivoli Access Manager authentication and authorization protection.

The JACC provider for Tivoli Access Manager is configured by default. The following list shows the JACC provider configuration settings for Tivoli Access Manager:

Field	Value
Name	Tivoli Access Manager
Description	This field is optional and used as a reference.
J2EE policy class name	com.tivoli.pd.as.jacc.TAMPolicy
Policy configuration factory class name	com.tivoli.pd.as.jacc.TAMPolicyConfigurationFactory
Role configuration factory class name	com.tivoli.pd.as.jacc.TAMRoleConfigurationFactory
JACC provider initialization class name	com.tivoli.pd.as.jacc.cfg.TAMConfigInitialize
Requires the EJB arguments policy context handler for access decisions	false
Supports dynamic module updates	true

To enable the JACC provider for Tivoli Access Manager, use the previous settings and complete the following steps:

1. Click **Security > Global security > External authorization providers**.
2. Select the **External authorization using a JACC provider** option, then click **Apply**.
3. Under Related Items, click **External JACC provider**. The JACC provider settings for Tivoli Access Manager are displayed.
4. Verify that the correct settings are present to work with your Tivoli Access Manager configuration. For more information, see “External Java Authorization Contract for Containers provider settings” on page 539.
5. Under Additional properties, click **Tivoli Access Manager properties**.
6. Click the **Enable embedded Tivoli Access Manager** option and verify that the correct Tivoli Access Manager server and WebSphere Application Server settings exist. For more information, see “Tivoli Access Manager JACC provider settings” on page 546.
7. Click **OK**.
8. Save the settings by clicking **Save** at the top of the page.
9. Log out of the WebSphere Application Server administrative console.
10. Restart WebSphere Application Server. The security configuration is now replicated to managed servers and node agents. These other servers within a cell also require restarting before the security changes take effect.

Enabling embedded Tivoli Access Manager

Embedded Tivoli Access Manager is not enabled by default, and you need to configure it for use.

About this task

Enabling Tivoli Access Manager security within WebSphere Application Server requires:

- A supported Lightweight Directory Access Protocol (LDAP) installed somewhere on your network. This user registry contains the user and group information for both Tivoli Access Manager and WebSphere Application Server.
- Tivoli Access Manager server exists and is configured to use the user registry. For details on the installation and configuration of Tivoli Access Manager, refer to the IBM Tivoli Access Manager for e-business information center.

Note: WebSphere Application Server contains an embedded client for Tivoli Access Manager. To use Tivoli Access Manager, you must also configure the Tivoli Access Manager server.

However, the server version must be the same version or later as the client version. For information on the supported version of Tivoli Access Manager, see WebSphere Application Server - Supported Prerequisites.

Tivoli Access Manager server is bundled with WebSphere Application Server.

- WebSphere Application Server is installed either in a single server model or as WebSphere Application Server Network Deployment.
- When administrative security is configured with a Federal Information Processing Standard (FIPS) provider, the Tivoli Access Manager server must be configured for FIPS as well

Complete the following steps to enable embedded Tivoli Access Manager security:

1. Create the security administrative user.
For more information, see the *Securing applications and their environment* PDF.
2. Configure the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager .
For more information, see the *Securing applications and their environment* PDF.
3. Enable WebSphere Application Server security. When you are using Tivoli Access Manager you must configure LDAP as the user registry.
For more information, see the *Securing applications and their environment* PDF.
4. Enable the JACC provider for Tivoli Access Manager.
For more information, see the *Securing applications and their environment* PDF.

Disabling embedded Tivoli Access Manager client using the administrative console

To unconfigure the JACC provider for Tivoli Access Manager, you can use the WebSphere Application Server administrative console.

About this task

Ensure that all the managed servers, including node agents, are started. Perform the following process once on the deployment management server. Information from the unconfigure operation is forwarded to managed servers, including node agents, when the server is restarted. The managed servers require their own restart for changes to take effect.

Note: It is also possible to unconfigure using the wsadmin command line utility. For details on unconfiguring the embedded Tivoli Access Manager client using the wsadmin command line utility, refer to “Disabling embedded Tivoli Access Manager client using wsadmin” on page 1059.

1. Restart the deployment manager process.

2. Click **Security > Global security > External authorization providers**.
3. Make sure that the default option, **Default authorization**, is checked, then click **OK**.
4. On the Global security panel, click **External authorization > External JACC provider**.
5. Under Additional properties, click **Tivoli Access Manager Properties**. The configuration screen for the JACC provider for Tivoli Access Manager is displayed.
6. Clear the **Enable embedded Tivoli Access Manager** option. If you want to ignore errors when unconfiguring, select the **Ignore errors during embedded Tivoli Access Manager disablement** option. Select this option only when the Tivoli Access Manager domain is in an irreparable state.
7. Click **OK**.
8. Optional: If you want security enabled without Tivoli Access Manager re-enable administrative security.
9. Optional: Synchronize all nodes.
10. Restart all WebSphere Application Server instances for the changes to take effect.

Forcing the unconfiguration of the Tivoli Access Manager JACC provider

If you find you cannot restart WebSphere Application Server after configuring the JACC provider for Tivoli Access Manager a utility is available to clear the security configuration and return WebSphere Application Server to an operable state.

About this task

The utility removes all of the PDLoginModuleWrapper entries as well as the Tivoli Access Manager authorization table from `security.xml` and `wsjaas.conf` files. This utility effectively removes the JACC provider for Tivoli Access Manager.

1. Back up the `security.xml` and `wsjaas.conf` files.
2. Enter the following command as one continuous line.

```
app_server_root/java/jre/bin/java
-classpath "app_server_root /$WAS_HOME/plugin/com.ibm.ws.runtime_1.0.0.jar"
com.tivoli.pd.as.jacc.cfg.CleanSecXML
fully_qualified_path/security.xml fully_qualified_path/wsjaas.conf
```

Propagating security policies and roles for previously deployed applications

Use this task to propagate security policies and roles to the external Java Authorization Contract for Containers (JACC) provider.

Before you begin

The external JACC provider must be configured before following these steps.

About this task

After switching to use the external JACC provider you can follow these steps to avoid having to redeploy your existing applications. Updating using these steps retrieves the security policy and roles from the deployed applications and propagates it to the external JACC provider removing the need for the applications to be redeployed.

1. From the WebSphere Application Server administrative console, click **Security → Global security → External authorization providers**.
2. Select the appropriate security policy and role updating option.
 - Select **Don't update provider** to not propagate any security policies or roles
 - Select **Update with all applications** to propagate security policies and roles for all applications
 - Select **Update with application names listed** to propagate security policies and roles for the selected applications. If multiple applications should be updated, separate the application names with commas.

3. Click **Apply**.

Results

After completing this task your security policies and roles have been successfully propagated to the external JACC provider.

Authorizing access to administrative roles

You can assign users and groups to administrative roles to identify users who can perform WebSphere Application Server administrative functions.

Before you begin

Administrative roles enable you to control access to WebSphere Application Server administrative functions. Refer to the descriptions of these roles in `rsec_adminroles.dita`.

- Before you assign users to administrative roles, you must set up your user registry. For information on the supported registry types, see “Selecting a registry or repository” on page 111.
- The following steps are needed to assign users to administrative roles.

About this task

You use the administrative console to assign users and groups to administrative roles and to identify users who can perform WebSphere Application Server administrative functions. In the administrative console,

1. Click **Users and Groups**. Click either **Administrative User Roles** or **Administrative Group Roles**.
2. To add a user or a group, click **Add** on the Console users or Console groups panel.
3. To add a new administrator user, follow the instructions on the page to specify a user, and select the **Administrator** role. Once the user is added to the Mapped to role list, click **OK**. The specified user is mapped to the security role.
4. To add a new administrative group, follow the instructions on the page to specify either a group name or a Special subject, highlight the **Administrator** role, and click **OK**. The specified group or special subject is mapped to the security role.
5. To remove a user or group assignment, click **Remove** on the Console Users or the Console Groups panel. On the Console Users or the Console Groups panel, select the check box of the user or group to remove and click **OK**.
6. To manage the set of users or groups to display, click **Show filter function** on the User Roles or Group Roles panel. In the **Search term(s)** box, type a value, then click **Go**. For example, `user*` displays only users with the user prefix.
7. After the modifications are complete, click **Save** to save the mappings.
8. Restart the application server for changes to take effect.
9. Shut down the nodes, node agents, and the deployment manager.
10. Verify that Java processes are not running. If they are running, discontinue these processes.
11. Restart the deployment manager.
12. Resynchronize the nodes. To resynchronize the nodes, run the `install_root/bin/syncNode` or the `install_root/bin/syncNode.sh` command for each node. For more information, see the `syncNode` command in the documentation.
13. Restart the nodes. To restart the nodes, run the `install_root/bin/startNode` or the `install_root/bin/startNode.sh` command for each node. For more information, see the `startNode` command in the documentation.
14. Start any clusters, if applicable.

What to do next

After you assign users to administrative roles, you must restart the Deployment Manager for the new roles to take effect. However, the administrative resources are not protected until you enable security.

Administrative user roles settings and CORBA naming service user settings

Use the Administrative User Roles page to give users specific authority to administer application servers through tools such as the administrative console or wsadmin scripting. The authority requirements are only effective when global security is enabled. Use the Common Object Request Broker Architecture (CORBA) naming service users settings page to manage CORBA naming service users settings.

To view the Console Users administrative console page, complete either of the following steps:

- Click **Security > Global security > Administrative User Roles**.
- Click **Users and Groups > Administrative User Roles**.

To view the CORBA naming service groups administrative console page, click **Environment > Naming > CORBA Naming Service Groups**.

User (Administrative user roles)

Specifies users.

The users that are entered must exist in the configured active user registry.

Data type: String

User (CORBA naming service users)

Specifies CORBA naming service users.

The users that are entered must exist in the configured active user registry.

Data type: String

Role (Administrative user roles)

Specifies user roles.

The following administrative roles provide different degrees of authority that are needed to perform certain application server administrative functions:

Administrator

The administrator role has operator permissions, configurator permissions, and the permission that is required to access sensitive data including server password, Lightweight Third Party Authentication (LTPA) password and keys, and so on.

Operator

The operator role has monitor permissions and can change the run-time state. For example, the operator can start or stop services.

Configurator

The configurator role has monitor permissions and can change the WebSphere Application Server configuration.

Deployer

The deployer role can complete both configuration actions and run-time operations on applications.

Monitor

The monitor role has the least permissions. This role primarily confines the user to viewing the application server configuration and current state.

adminsecuritymanager

The adminsecuritymanager role has privileges for managing users and groups from within the administrative console and determines who has access to modify users and groups using administrative role mapping. Only the adminsecuritymanager role can map users and groups to administrative roles, and by default, AdminId is granted to the adminsecuritymanager.

iscadmins

The iscadmins role has administrator privileges for managing users and groups from within the administrative console only.

Note: To manage users and groups, click **Users and Groups** in the console navigation tree. Click either **Manage Users** or **Manage Groups**.

Data type: String
Range: Administrator, Operator, Configurator, Deployer, Monitor, and iscadmins

Note: Other arbitrary administrative roles might also be visible in the administrative console collection table. Other contributors to the console might create these additional roles, which can be used for applications that are deployed to the console.

Role (CORBA naming service users)

Specifies naming service user roles.

A number of naming roles are defined to provide degrees of authority that are needed to perform certain application server naming service functions. The authorization policy is only enforced when global security is enabled. The following roles are valid: CosNamingRead, CosNamingWrite, CosNamingCreate, and CosNamingDelete.

The roles now have authority levels from low to high:

CosNamingRead

You can query the application server name space by using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The EVERYONE special-subject is the default policy for this role.

CosNamingWrite

You can perform write operations such as JNDI bind, rebind, or unbind, plus CosNamingRead operations.

CosNamingCreate

You can create new objects in the name space through operations such as JNDI createSubcontext and CosNamingWrite operations.

CosNamingDelete

You can destroy objects in the name space, for example using the JNDI destroySubcontext method and CosNamingCreate operations.

Data type: String
Range: CosNamingRead, CosNamingWrite, CosNamingCreate and CosNamingDelete

Login status (Administrative user roles)

Specifies whether the user is active or inactive.

Administrative group roles and CORBA naming service groups

Use the Administrative Group Roles page to give groups specific authority to administer application servers through tools such as the administrative console or wsadmin scripting. The authority requirements are only

effective when administrative security is enabled. Use the Common Object Request Broker Architecture (CORBA) naming service groups page to manage CORBA Naming Service groups settings.

To view the Console Groups administrative console page, complete either of the following steps:

- Click **Security > Global security > Administrative Group Roles**.
- Click **Users and Groups > Administrative Group Roles**.

To view the CORBA naming service groups administrative console page, click **Environment > Naming > CORBA Naming Service Groups**.

Group (CORBA naming service groups)

Identifies CORBA naming service groups.

In previous releases of WebSphere Application Server, there were two default groups: ALL AUTHENTICATED and EVERYONE. However, EVERYONE is now the only default group, and it provides CosNamingRead privileges only.

Data type: String
Range: EVERYONE

Role (CORBA naming service groups)

Identifies naming service group roles.

A number of naming roles are defined to provide the degrees of authority that are needed to perform certain application server naming service functions. The authorization policy is only enforced when global security is enabled.

Four name space security roles are available: CosNamingRead, CosNamingWrite, CosNamingCreate, and CosNamingDelete. The roles have authority levels from low to high:

Cos Naming Read

You can query the application server name space using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The EVERYONE special-subject is the default policy for this role.

Cos Naming Write

You can perform write operations such as JNDI bind, rebind, or unbind, and CosNamingRead operations. The ALL_AUTHENTICATED special-subject is the default policy for this role.

Cos Naming Create

You can create new objects in the name space through operations such as JNDI createSubcontext and CosNamingWrite operations. The ALL_AUTHENTICATED special-subject is the default policy for this role.

Cos Naming Delete

You can destroy objects in the name space, for example using the JNDI destroySubcontext method and CosNamingCreate operations. The ALL_AUTHENTICATED special-subject is the default policy for this role.

Data type: String
Range: CosNamingRead, CosNamingWrite, CosNamingCreate, and CosNamingDelete

Group (Administrative group roles)

Specifies groups.

The ALL_AUTHENTICATED and the EVERYONE groups can have the following role privileges: Administrator, Configurator, Operator, and Monitor.

Data type: String
Range: ALL_AUTHENTICATED, EVERYONE

Role (Administrative group roles)

Specifies user roles.

The following administrative roles provide different degrees of authority needed to perform certain application server administrative functions:

Administrator

The administrator role has operator permissions, configurator permissions, and the permission that is required to access sensitive data, including server password, Lightweight Third Party Authentication (LTPA) password and keys, and so on.

Operator

The operator role has monitor permissions and can change the run-time state. For example, the operator can start or stop services.

Configurator

The configurator role has monitor permissions and can change the application server configuration.

Deployer

The deployer role can perform both configuration actions and runtime operations on applications.

Monitor

The monitor role has the least permissions. This role primarily confines the user to viewing the application server configuration and current state.

iscadmins

The iscadmins role has administrator privileges for managing users and groups from within the administrative console only.

Note: To manage users and groups, click **Users and Groups** in the console navigation tree. Click either **Manage Users** or **Manage Groups**.

Auditor

The auditor can view and modify the configuration settings for the security auditing subsystem. The auditor role includes the monitor role.

Data type: String
Range: Administrator, Operator, Configurator, Monitor, Deployer and iscadmins

Note: Other arbitrary administrative roles might also be visible in the administrative console collection table. Other contributors to the console might create these additional roles, which can be used for applications that are deployed to the console.

Assigning users to naming roles

Use this task to assign users to naming roles by using the administrative console.

About this task

The following steps are needed to assign users to naming roles. In the administrative console, click **Environment > Naming**, and click **CORBA Naming Service Users** or **CORBA Naming Service Groups**.

1. Click **Add** on the CORBA Naming Service Users or the CORBA Naming Service Groups panel.

2. To add a new naming service user, follow the instructions on the page to specify a user, and select one or more roles. Once the user is added to the "Mapped to role" list, click **OK**. The specified user is mapped to one or more security roles.
3. To add a new naming service group, follow the instructions on the page to specify either a group name or a Special subject, highlight one or more roles, and click **OK**. The specified group or special subject are mapped to one or more the security roles
4. To remove a user or group assignment, go to the **CORBA Naming Service Users** or **CORBA Naming Service Groups** panel. Select the check box next to the user or group that you want to remove and click **Remove**.
5. To manage the set of users or groups to display, expand the **Filter** folder on the right panel, and modify the filter text box. For example, setting the filter to user* displays only users with the user prefix.
6. After modifications are complete, click **Save** to save the mappings. Restart the server for the changes to take effect.

Example

The default naming security policy is to grant all users read access to the CosNaming space and to grant any valid user the privilege to modify the contents of the CosNaming space. You can perform the previously mentioned steps to restrict user access to the CosNaming space. However, use caution when changing the naming security policy. Unless a Java Platform, Enterprise Edition (Java EE) application has clearly specified its naming space access requirements, changing the default policy can result in unexpected org.omg.CORBA.NO_PERMISSION exceptions at runtime.

Propagating administrative role changes to Tivoli Access Manager

These steps provide an example of how to migrate the admin-authz.xml file.

About this task

Additions and changes to console users and groups are not automatically added to the Tivoli Access Manager object space after the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager is configured. Changes to console users and groups are saved in the admin-authz.xml file and this file must be migrated before any changes take effect. The JACC provider for Tivoli Access Manager includes the **migrateEAR** migration utility for incorporating console user and group changes into the Tivoli Access Manager object space.

Note: The **migrateEAR** utility is used to migrate the changes made to console users and groups after the JACC provider for Tivoli Access Manager is configured. The utility does not need to run for changes and additions to console users and groups made prior to the configuration of the JACC provider for Tivoli Access Manager because the changes made to the admin-authz.xml and naming-authz.xml files are automatically migrated at configuration time. Furthermore, the migration tool does not need to run before deploying standard Java Platform, Enterprise Edition (Java EE) applications; Java EE application policy deployment is also performed automatically.

For example, if you wanted to migrate the admin-authz.xml file, perform the following steps:

1. Set up the environment.
Before running the migrateEAR utility, set up the environment by running the setupCmdLine.bat or setupCmdLine.sh file that is located in the *app_server_root/bin* directory.
Make sure that the *WAS_HOME* environment variable is set to the WebSphere Application Server installation directory.
2. Change to the *app_server_root/bin* directory where the migrateEAR utility is located.
3. Run the **migrateEAR** utility to migrate the data contained in the admin-authz.xml file. Use the parameter descriptions that are listed in "migrateEAR utility for Tivoli Access Manager" on page 567.

UNIX

For example:

```
migrateEAR
-j "app_server_root/profiles/profile_name/config/cells/cell_name/xml_filename"
-a sec_master
-p password
-w wsadmin
-d o=ibm,c=us
-c file:"app_server_root/java/jre/PdPerm.properties"
-z Roles
```

where *xml_filename* might be *admin-authz.xml* or *naming-authz.xml*. The `-z Roles` parameter is optional and when specified adds a subdirectory under the current directory structure in which to store the role mapping. For example,

```
/WebAppServer/deploYedResouces/Roles
```

If `-z Roles` is not specified, the role mapping is stored in the current directory structure. For example,

```
/WebAppServer/deploYedResouces
```

A status message is displayed when the migration completes. Output of the utility is logged to the `pdwas_migrate.log` file, which is created in the directory where the utility is run. Check the log file after each migration. If the log file displays errors, check the last recorded transaction, correct the source of the error, and rerun the migration utility. If the migration is unsuccessful, verify that you supplied the correct values for the `-c` and `-j` options.

4. WebSphere Application Server does not require a restart for the changes to take effect.

migrateEAR utility for Tivoli Access Manager

The **migrateEAR** utility migrates changes made to console users and groups in the `admin-authz.xml` and `naming-authz.xml` files into the Tivoli Access Manager object space.

Syntax

```
migrateEAR
-j fully_qualified_filename
-c pdPerm.properties_file_location
-a Tivoli_Access_Manager_administrator_ID
-p Tivoli_Access_Manager_administrator_password
-w WebSphere_Application_Server_administrator_user_name
-d user_registry_domain_suffix
[-r root_objectspace_name]
[-t ssl_timeout]
[-z role_mapping_location]
```

Parameters

Note: In the following parameters, use the absolute path instead of a variable.

-a *Tivoli_Access_Manager_administrator_ID*

The administrative user identifier. The administrative user must have the privileges required to create users, objects, and access control lists (ACLs). For example, `-a sec_master`.

This parameter is optional. When the parameter is not specified, you are prompted to supply it at run time.

-c *PdPerm.properties_file_location*

The Uniform Resource Indicator (URI) location of the `PdPerm.properties` file that is configured by the `pdwascfg` utility. When WebSphere Application Server is installed in the default location, the URI is:

Solaris

Linux

HP-UX

```
file:/opt/IBM/WebSphere/AppServer/java/jre/PdPerm.properties
```

AIX

```
file:/usr/IBM/WebSphere/AppServer/java/jre/PdPerm.properties
```

Windows

```
file:"C:/Program Files/IBM/WebSphere/AppServer/java/jre/PdPerm.properties"
```

-d *user_registry_domain_suffix*

The domain suffix for the user registry to use. For example, for Lightweight Directory Access Protocol (LDAP) user registries, this value is the domain suffix, such as: "o=ibm,c=us"

Windows

Windows platforms require that the domain suffix is enclosed within quotes.

You can use the **pdadmin user show** command to display the distinguished name (DN) for a user.

-j *fully_qualified_pathname*

The fully qualified path and file name of the Java 2 Platform, Enterprise Edition application archive file ,admin-authz.xml or the roles definitions file naming-authz.xml that is used for a naming operation authorization. Optionally, this path can also be a directory of an expanded enterprise application. For example, when WebSphere Application Server is installed in the default location, the path to the data files to migrate includes:

Solaris

Linux

HP-UX

```
file:/opt/IBM/WebSphere/AppServer/profiles/profile_name/config/cells  
/cell_name/admin-authz.xml
```

AIX

```
file:/usr/IBM/WebSphere/AppServer/profiles/profile_name/config/cells  
/cell_name/admin-authz.xml
```

Windows

```
"C:/Program Files/IBM/WebSphere/AppServer/profiles/profile_name/config/cells  
/cell_name/admin-authz.xml"
```

-p *Tivoli_Access_Manager_administrator_password*

The password for the Tivoli Access Manager administrative user. The administrative user must have the privileges that are required to create users, objects, and access control lists (ACLs). For example, you can specify the password for the -a sec_master administrative user as -p myPassword.

When this parameter is not specified, the user is prompted to supply the password for the administrative user name.

-r *root_objectspace_name*

The space name of the root object. The value is the name of the root of the protected object namespace hierarchy that is created for WebSphere Application Server policy data.

The default value for the root object space is WebAppServer.

Set the Tivoli Access Manager root object space name by modifying the amwas.amjacc.template.properties file prior to configuring the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager for the first time. Use this option if the default object space value is not used in the configuration of the Tivoli Access Manager JACC provider for Tivoli Access Manager.

Do not change the Tivoli Access Manager object space name after the Tivoli Access Manager JACC provider is configured.

-t *ssl_timeout*

The number of minutes for the Secure Sockets Layer (SSL) timeout. This parameter is used to disconnect and reconnect the SSL context between the Tivoli Access Manager authorization server and the policy server before the default connection times out.

The default is 60 minutes. The minimum value is 10 minutes. The maximum value cannot exceed the Tivoli Access Manager `ssl-v3-timeout` value. The default value for `ssl-v3-timeout` is 120 minutes.

If you are not familiar with the administration of this value, you can safely use the default value.

-w *WebSphere_Application_Server_administrator_user_name*

The user name that is configured in the WebSphere Application Server security user registry field as the administrator. This value matches the account that you created or imported in “Creating the security administrative user” on page 544. Access permission for this user is needed to create or update the Tivoli Access Manager protected object space.

When the WebSphere Application Server administrative user does not already exist in the protected object space, it is created or imported. In this case, a random password is generated for the user and the account is set to `not valid`. Change this password to a known value and set the account to `valid`.

A protected object and access control list (ACL) are created. The administrative user is added to the `pdwas-admin` group with the following ACL attributes:

T Traverse permission
i Invoke permission

WebAppServer

You can overwrite the action group name. The default name is `WebAppServer`. This action group name and the matching root object space can be overwritten when the migration utility is run with the `-r` option.

-z *role_mapping_location*

The location where the role mapping is to be stored when migrating administration applications. The default location is to place the role mapping in the current directory structure, such as:

```
/WebAppServer/deploYedResouces
```

Specifying the `-z` option adds another directory level in which to store the role mapping. For example, if you specify `-z Roles` in the `migrateEAR` utility, the role mapping is stored in the directory structure as follows:

```
/WebAppServer/deploYedResouces/Roles
```

Comments

This utility migrates security policy information from deployment descriptors or enterprise archive files to Tivoli Access Manager for WebSphere Application Server. The script calls `com.tivoli.pdwas.migrate.Migrate` the Java class.

Before invoking the script you must run the **setupCmdLine.bat** or the **setupCmdLine.sh** commands. These files can be found in the `%WAS_HOME%/bin` directory.

The script is dependent on finding the correct environment variables for the location of prerequisite software.

The script calls Java code with the following options:

-Dpdwas.lang.home

The directory that contains the native language support libraries that are provided with the JACC provider for Tivoli Access Manager. These libraries are located in a subdirectory under the JACC provider for Tivoli Access Manager installation directory. For example: `-Dpdwas.lang.home=%PDWAS_HOME%\java\nls`

-cp %CLASSPATH% com.tivoli.pdwas.migrate.Migrate

The `CLASSPATH` variable must be set correctly for your Java installation.

Windows Both the -j option and the -c option can reference the %WAS_HOME% variable to determine where WebSphere Application Server is installed. This information is used to:

- Build the full path name of the enterprise archive file.
- Build the full URI path name to the location of the PdPerm.properties file.

To enable a new user access to the administrative group in WebSphere Application Server, it is recommended that the user be added to the pdwas-admin group after JACC has been enabled. You can enter the administrative primary ID (adminID) in the group. This is required when the serverID is not the same as the adminID.

The following is an example of this command:

```
pdadmin> group modify pdwas-admin add adminID
```

Return codes

The utility can return the following exit status codes:

- 0 The command completed successfully.
- 1 The command failed.

Assigning users from a foreign realm to the admin-authz.xml

Operating with the administrative agent and job manager topology allows more situations where you might need to add an administrative user from a different registry into your administrative authorization table (admin-authz.xml). Each administrative user that needs to be added requires the "accessID" format of the user from the remote registry. When that user finally is active in the local cell, the authorization table will already have that accessID that is required. This task demonstrates how this assignment of users is performed.

1. You need to determine the accessId for a user on the remote registry. To do this, you call the following wsadmin task and query based on a user filter. The following example illustrates a query from the registry realm "BIRKT60" with a userFilter of "localuser*". This query returns any user from this realm that begins with "localuser". The resulting accessId is the one you need to specify in the target administrative authorization table in the following step. Connect to the sending administrative process:

```
wsadmin> $AdminTask listRegistryUsers {-securityRealmName BIRKT60 -displayAccessIds true -userFilter localuser*}  
{name BIRKT60\localuser@BIRKT60}  
{accessId user:BIRKT60/S-1-5-21-3033296400-14683092-2821094880-1007}
```

2. Add "localuser" to the target admin-authz.xml using the following wsadmin task. Connected to the receiving administrative process:

```
wsadmin> $AdminTask mapUsersToAdminRole {-roleName administrator -userids {localuser } -accessids {user:BIRKT60/S-1-5-2
```

3. Save the changes.

Results

This task updates the admin-authz.xml in the receiving administrative process to allow a "cross-realm authorization" to succeed. The example illustrated here was for a LocalOS registry user. Performing the same task for an LDAP accessId produces results that look more like a realm and distinguished name (DN).

Fine-grained administrative security

In releases prior to WebSphere Application Server version 6.1, users granted administrative roles could administer all of the resource instances under the cell. WebSphere Application Server is now more fine-grained, meaning that access can be granted to each user per resource instance.

For example, users can be granted configurator access to a specific instance of a resource only (an application, an application server or a node). Users cannot access any other resources outside of the resources assigned to them. The administrative roles are now per resource instance rather than to the entire cell. However, there is a cell-wide authorization group for backward compatibility. Users assigned to administrative roles in the cell-wide authorization group can still access all of the resources within the cell.

Note: Nodes prior to WebSphere Application Server Version 6.1 in a mixed cell environment are filtered out of resource mapping.

To achieve this instance-based security or fine-grained security, resources that require the same privileges are placed in a group called the *administrative authorization group* or *authorization group*. Users can be granted access to the authorization group by assigning to them the required administrative role.

Fine-grained administrative security can also be used in single-server environments. Various applications in the single server can be grouped and placed in different authorization groups. Therefore, there are different authorization constraints for different applications. Note that the server itself cannot be part of any authorization group in a single-server environment.

You can assign users and groups to the adminsecuritymanager role on the cell level through wsadmin scripts and the administrative console. Using the adminsecuritymanager role, you can assign users and groups to the administrative user roles and administrative group roles.

When fine grained administrative security is used, users granted the adminsecuritymanager role can manage authorization groups. See “Administrative roles and naming service authorization” on page 510 for detailed explanations of all administrative roles.

An administrator cannot assign users and groups to the administrative user roles and administrative group roles, including the adminsecuritymanager role. See “Administrative roles” on page 518 for more details.

There are several administrative security commands that can be used to create authorization groups, map resources to authorization groups, and to assign users to administrative roles within the authorization groups. Following are some examples using wsadmin:

- **Create a new authorization group:**

```
$AdminTask createAuthorizationGroup {-authorizationGroupName authGroup1}
```
- **Deleting an authorization group:**

```
$AdminTask deleteAuthorizationGroup {-authorizationGroupName groupName}
```
- **Add resources to an authorization group:**

```
$AdminTask addResourceToAuthorizationGroup  
{-authorizationGroupName groupName -resourceName Application=appl}
```
- **Remove resources from an authorization group:**

```
$AdminTask removeResourceFromAuthorizationGroup  
{-authorizationGroupName groupName -resourceName Application=appl}
```
- **Add user IDs to roles in an authorization group:**

```
$AdminTask mapUsersToAdminRole {-authorizationGroupName groupName  
-roleName administrator -userids user1}
```
- **Add group IDs to roles in an authorization group:**

```
$AdminTask mapGroupsToAdminRole {-authorizationGroupName groupName  
-roleName administrator -groupids group1}
```
- **Remove user IDs from roles in an authorization group:**

```
AdminTask removeUsersFromAdminRole {-authorizationGroupName  
groupName -roleName administrator -userids user1}
```
- **Remove group IDs from roles in an authorization group:**

```
$AdminTask removeGroupsFromAdminRole {-authorizationGroupName  
groupName -roleName administrator -groupids group1}
```


Resources that can be added to an authorization group


You can add only resource instances of the following types to an authorization group:

- Cell
- Node
- ServerCluster
- Server
- Application
- NodeGroup

If a resource instance is not one of the types listed above, its parent resource will be used.

A resource instance can only belong to one authorization group. However, there is a containment relationship among resource instances. If a parent resource belongs to a different authorization group than that of its child resource instance, the child resource instance implicitly will belong to multiple authorization groups. You cannot add the same resource instance to more than one authorization group.

The following diagram shows the containment relationship among resource instances:



The privileges required for actions on resource instances depend on two factors:

- The authorization group of the administrative resource instance. If a user is granted access to an authorization group, all of the resource instances in that group will be included.
- The containment relationship of the resource instance. If a user is granted access to a parent resource instance, all of the children resource instances will be included.

The privileges required to access various administrative resource instances are shown in the following table:

Resource	Action	Required roles
Server	Start, stop, runtime operations	Server-operator, node-operator, cell-operator
Server	New, delete	Node-configurator, cell-configurator
Server	Edit configuration	Server-configurator, node-configurator, cell-configurator
Server	View configuration, runtime status	Server-monitor, node-monitor, cell-monitor
Node	Restart, stop, sync	Node-operator, Cell-operator
Node	Add, delete	Cell-configurator
Node	Edit configuration	Node-configurator, cell-configurator

Resource	Action	Required roles
Node	View configuration, runtime status	Node-monitor, cell-monitor
Cluster	Start, stop, runtime operations	Cluster-operator, cell-operator
Cluster	New, delete	Cell-configurator
Cluster	Edit configuration	Cluster-configurator, cell-configurator
Cluster	View configuration, runtime status	Cluster-monitor, cell-monitor
Cluster member	Start, stop, runtime operations	Server-operator, cluster-operator, node-operator, cell-operator
Cluster member	New, delete	Node-configurator, cell-configurator
Cluster member	Edit configuration	Server-configurator, cluster-configurator, node-configurator, cell-configurator
Cluster member	View configuration, runtime status	Server-monitor, cluster-monitor, node-monitor, cell-monitor
Application	All operations	Refer to the section "Deployer roles" in "Administrative roles" on page 518.
Node, cluster	Add, delete	Cell-configurator

The *server-operator* role is the operator role of the authorization group to which the server instance is part of. Similarly, the *node-operator* role is in the operator role of the authorization group to which the node instance is part of.

To use fine-grained administrative security in the administrative console, a user should be granted a monitor role at the cell level at minimum. However, to login using wsadmin, a user should be granted a monitor role for any authorization group.

If you log in to the administrative console as a cell-level administrator, adminsecuritymanager, operator, monitor or configurator, you can perform all of the operations that the role allows you to do. However, if you want to give users access only to specific authorization groups or permissions to non-cell authorizations groups, you must use **wsadmin**.

New Administrative Authorization Group

Use this page to create a new administrative authorization group and to specify the associated administrative resources.

To view this administrative console page, click **Security** → **Administrative Authorization Groups** → **New**.

You must be logged into the administrative console with cell-level **AdminSecurityManager** authority, or the primary administrative ID can make these changes as well.

Related concepts

“Role-based authorization” on page 516

Use authorization information to determine whether a caller has the necessary privileges to request a service.

“Administrative roles and naming service authorization” on page 510

WebSphere Application Server extends the Java Platform, Enterprise Edition (Java EE) security role-based access control to protect the product administrative and naming subsystems.

Related reference

“Example: Using fine-grained security” on page 581

The following scenarios describe the use of fine-grained administrative security, particularly the new deployment role.

Name

Use to identify the new administrative authorization group. The name should be descriptive of the group’s role, or purpose, and should be unique in the cell structure. Using a non-unique name results in an error and a failure to create the new administrative authorization group. This is a required field.

Resources

Select the resources from the Resource section to which you want the new administrative authorization group to control access.

Resources that are displayed in black text are available for selection.

Resources that are displayed in grey are already members of a different administrative authorization group. Therefore, these resources are not available for inclusion in the new administrative authorization group. When a resource is a member of a different authorization group, the name of the group displays next to the resource name. For example: server_1 (group_1)

Your filtering options include the following:

Nodes

All of the resources associated with the node.

Clusters

All of the resources associated with the cluster.

Applications

All of the resources associated with the application.

Web Servers

All of the resources associated with the Web servers.

Servers

All of the resources associated with the servers.

Node Groups

All of the resources associated with the Node Groups.

All resources

The default view that displays the authorization group tree.

Assigned resources

Displays all of the resources explicitly assigned to the current authorization group.

Administrative Authorization Group collection

Use this page to create, delete or to edit an existing administrative authorization group.

To view this administrative console page, click **Security** → **Administrative Authorization Groups**.

You must be logged into the administrative console with cell-level **AdminSecurityManager** authority, or the primary administrative ID can make these changes as well.

Related concepts

“Role-based authorization” on page 516

Use authorization information to determine whether a caller has the necessary privileges to request a service.

“Administrative roles and naming service authorization” on page 510

WebSphere Application Server extends the Java Platform, Enterprise Edition (Java EE) security role-based access control to protect the product administrative and naming subsystems.

Related reference

“Example: Using fine-grained security” on page 581

The following scenarios describe the use of fine-grained administrative security, particularly the new deployment role.

Name

The name field specifies the current name of the administrative authorization group. You can edit the name of the administrative authorization group during the creation process only. After the authorization group is created, you cannot modify the name. The specified name must be unique within the cell structure. Otherwise, a non-unique name results in an error.

New

Select to create a new administrative authorization group.

Delete

Select to remove an existing administrative authorization group.

Note: You must select an administrative authorization group before selecting **Delete**.

Creating a fine-grained administrative authorization group using the administrative console

You can create a fine-grained administrative authorization group by selecting administrative resources to be part of the authorization group. You can assign users or groups to this new administrative authorization group and also give them access to the administrative resources contained within.

Before you begin

You must be logged into the administrative console with cell-level **Admin Security Manager** authority, or the primary administrative ID can make these changes as well.

1. Navigate to **Security** → **Administrative Authorization Groups** → **New**.
2. Type a name for the administrative authorization group into the **Name** field. This is a required field. The name must be unique within the cell structure. If the name is not unique then the new administrative authorization group is not created at the end of this procedure.
3. Select the resources from the Resource section to which you want the new administrative authorization group to control access.

Resources that are displayed in black text are available for selection.

Resources that are displayed in grey are already members of a different administrative authorization group. Therefore, these resources are not available for inclusion in the new administrative authorization group. When a resource is a member of a different authorization group, the name of the group displays next to the resource name. For example: server_1 (group_1)

Your filtering options include the following:

- Nodes. (All of the resources associated with the nodes.)

- Servers. (All of the resources associated with the servers.)
 - Web servers. (All of the resources associated with the Web servers.)
 - Clusters. (All of the resources associated with the clusters.)
 - Applications. (All of the resources associated with the applications.)
 - Node groups. (All of the resources associated with the Node Groups.)
 - All scopes. (The default view that displays the authorization group tree.)
 - Assigned scopes. (Displays all of the scopes explicitly assigned to the current authorization group.)
4. Click **OK** or **Apply**.
 5. If you want to associate a user role to this new administrative authorization group, do the following:
 - a. Click **Administrative user roles** located under the Additional Properties section. The available user roles are the following:

Administrator

An individual or group that uses the administrator role has the operator and configurator privileges plus additional privileges that are granted solely to the administrator role. For example, an administrator can complete the following tasks:

- Modify the server user ID and password.
- Configure authentication and authorization mechanisms.
- Enable or disable administrative security.
- Enable or disable Java 2 security.
- Change the Lightweight Third Party Authentication (LTPA) password and generate keys.
- Create, update, or delete users in the federated repositories configuration.
- Create, update, or delete groups in the federated repositories configuration.

Note: An administrator cannot map users and groups to the administrator roles.

Configurator

An individual or group that uses the configurator role has the monitor privilege plus the ability to change the WebSphere Application Server configuration. The configurator can perform all the day-to-day configuration tasks. For example, a configurator can complete the following tasks:

- Create a resource.
- Map an application server.
- Install and uninstall an application.
- Deploy an application.
- Assign users and groups-to-role mapping for applications.
- Set up Java 2 security permissions for applications.
- Customize the Common Secure Interoperability Version 2 (CSlv2), Secure Authentication Service (SAS), and Secure Sockets Layer (SSL) configurations.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Deployer

Users granted this role can perform both configuration actions and runtime operations on applications..

Operator

An individual or group that uses the operator role has monitor privileges plus ability to change the runtime state. For example, an operator can complete the following tasks:

- Stop and start the server.

- Monitor the server status in the administrative console.

Monitor

An individual or group that uses the monitor role has the least amount of privileges. A monitor can complete the following tasks:

- View the WebSphere Application Server configuration.
- View the current state of the Application Server.

Admin Security Manager

Using the Admin Security Manager role, you can assign users and groups to the administrative user roles and administrative group roles. However, an administrator cannot assign users and groups to the administrative user roles and administrative group roles including the Admin Security Manager roles.

- b. Click **Add...** The New User page is displayed.
 - c. Select the appropriate role(s) from the **Role(s)** list box.
 - d. Select a user or users by entering text in the **Search string** field, and then click **Search..** Click the arrow to add the available user or users to the **Mapped to role** field. You can select multiple users and roles by clicking **Select All**.
 - e. Click **OK**. You are returned to the Administrative User Roles page. The new users are displayed in the Administrative User Roles table along with their appropriate roles.
 - f. Repeat steps B through E for each new user to whom you want to map a role.
6. If you want to associate a group to this new user role, do the following:
 - a. Click **Administrative group roles** located under the Additional Properties section.
 - b. Click **Add...** The New Group page is displayed.
 - c. Select the appropriate role or roles from the **Role(s)** list box.
 - d. Select a user or users by entering text in the **Search string** field, and then click **Search..** Click the arrow to add the available user or users to the **Mapped to role** field. You can select multiple users and roles by clicking **Select All**.
 - e. Select either the **Select from special subjects** or **Map Groups As Specified Below** option.

If you select the **Select from special subjects** option, you can select the EVERYONE, ALL AUTHENTICATED, or ALL AUTHENTICATED IN TRUSTED REALMS values.

A list of user groups and roles are displayed in the **Available** and **Mapped to role** fields. Select the user groups from the **Available** field and then select the roles from the **Mapped to role** field to which you want the group or groups associated. You can select multiple groups and roles.
 - f. Click **OK**. You are returned to the Administrative Group Roles page. The new group is displayed in the Administrative Group Roles table along with the role of the new group.
 - g. Repeat steps B through E for each new group to whom you want to map a role.
 7. If you want to create another administrative authorization group, click **Apply**. The current administrative authorization group is created. Repeat steps 2 through 6 to create another administrative authorization group.
 8. If you do not want to create another administrative authorization group, click **OK**.

Related concepts

“Role-based authorization” on page 516

Use authorization information to determine whether a caller has the necessary privileges to request a service.

“Administrative roles and naming service authorization” on page 510

WebSphere Application Server extends the Java Platform, Enterprise Edition (Java EE) security role-based access control to protect the product administrative and naming subsystems.

Related tasks

“Editing a fine-grained administrative authorization group using the administrative console”

You can add or remove administrative resources to an administrative authorization group or edit an existing one.

Editing a fine-grained administrative authorization group using the administrative console

You can add or remove administrative resources to an administrative authorization group or edit an existing one.

Before you begin

You must be logged into the administrative console with the cell-level **AdminSecurityManager** authority or as the primary administrative user.

1. Navigate to **Security** → **Administrative Authorization Groups**. The Administrative Authorization Groups page displays a table that lists all of the current administrative authorization groups available in the cell.
2. Click on the administrative authorization group in the table that you want to edit.
3. To add or remove resources from the administrative authorization group, select or clear them in the Resource section of the edit page. Resources displayed in black text are available for selection or clearing. Resources displayed in grey text are members of a different administrative authorization group and therefore cannot be edited for the current administrative authorization group.

Your filtering options include the following:

- Nodes. (All of the resources associated with the nodes.)
- Servers. (All of the resources associated with the servers.)
- Web servers. (All of the resources associated with the Web servers.)
- Clusters. (All of the resources associated with the clusters.)
- Applications. (All of the resources associated with the applications.)
- Node groups. (All of the resources associated with the Node Groups.)
- All scopes. (The default view that displays the authorization group tree.)
- Assigned scopes. (Displays all of the scopes explicitly assigned to the current authorization group).

Note: Nodes prior to WebSphere Application Server Version 6.1 in a mixed cell environment are filtered out of resource mapping.

4. To remove a user or a group, do the following:
 - a. To delete users, click **Administrative user roles** under the Additional Properties section. To delete groups, click **Administrative group roles** under the Additional Properties section. The appropriate edit page displays a table that lists all of the current users or groups and their associated roles, along with the user’s login status.
 - b. Click the check box beside the name of the current user or group and then click **Remove**. The current user or group is no longer associated with the role and the role is no longer listed in the table. It is now ready to have a new user or group assigned to it.

5. If you want to add or to reassign a user or group role to this administrative authorization group, do the following:
 - a. To add a user, click **Administrative user roles** under the Additional Properties section. To add a group, click **Administrative group roles** located under the Additional Properties section. The appropriate edit page displays a table that lists all of the current users or groups and their associated roles. The available roles are:

Administrator

An individual or group that uses the administrator role has the operator and configurator privileges plus additional privileges that are granted solely to the administrator role. For example, an administrator can complete the following tasks:

- Modify the server user ID and password.
- Configure authentication and authorization mechanisms.
- Enable or disable administrative security.
- Enable or disable Java 2 security.
- Change the Lightweight Third Party Authentication (LTPA) password and generate keys.
- Create, update, or delete users in the federated repositories configuration.
- Create, update, or delete groups in the federated repositories configuration.

Note: An administrator cannot map users and groups to the administrator roles.

Configurator

An individual or group that uses the configurator role has the monitor privilege plus the ability to change the WebSphere Application Server configuration. The configurator can perform all the day-to-day configuration tasks. For example, a configurator can complete the following tasks:

- Create a resource.
- Map an application server.
- Install and uninstall an application.
- Deploy an application.
- Assign users and groups-to-role mapping for applications.
- Set up Java 2 security permissions for applications.
- Customize the Common Secure Interoperability Version 2 (CSlv2), Security Authentication Service (SAS), and Secure Sockets Layer (SSL) configurations.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Deployer

Users granted this role can perform both configuration actions and runtime operations on applications.

Operator

An individual or group that uses the operator role has monitor privileges plus ability to change the runtime state. For example, an operator can complete the following tasks:

- Stop and start the server.
- Monitor the server status in the administrative console.

Monitor

An individual or group that uses the monitor role has the least amount of privileges. A monitor can complete the following tasks:

- View the WebSphere Application Server configuration.

- View the current state of the Application Server.

Admin Security Manager

Using the Admin Security Manager role, you can assign users and groups to the administrative user roles and administrative group roles. However, an administrator cannot assign users and groups to the administrative user roles and administrative group roles including the Admin Security Manager role.

- b. Click **Add...**
- c. To add a new user or group, follow the instructions on the page to specify either a user name, group name, or Special subject. Highlight the desired role(s), and click **OK**. The specified users, groups, or Special subject are mapped to the security roles.

Related concepts

“Role-based authorization” on page 516

Use authorization information to determine whether a caller has the necessary privileges to request a service.

“Administrative roles and naming service authorization” on page 510

WebSphere Application Server extends the Java Platform, Enterprise Edition (Java EE) security role-based access control to protect the product administrative and naming subsystems.

Related tasks

“Creating a fine-grained administrative authorization group using the administrative console” on page 575

You can create a fine-grained administrative authorization group by selecting administrative resources to be part of the authorization group. You can assign users or groups to this new administrative authorization group and also give them access to the administrative resources contained within.

Fine-grained administrative security in heterogeneous and single-server environments

Fine-grained administrative security can be used in heterogeneous or single-server environments with some restrictions.

Fine-grained administrative security in a heterogeneous environment

In WebSphere Application Server Version 6.0, heterogeneous systems are supported. Specifically, a deployment manager node can run in WebSphere Application Server Version 6.0, some nodes can run WebSphere Application Server Version 6.0, and other nodes can run WebSphere Application Server Version 5.x. In WebSphere Application Server Version 7.0, nodes are available for WebSphere Application Server Versions 5.x, 6.0, and Version 7.0.

Because all of the configurations that are done in deployment manager node are always of WebSphere Application Server Version 7.0 or higher, fine-grained administrative security can be enforced when configuring resources that belong to earlier releases. However, run-time code for versions lower than Version 7.0 cannot enforce fine-grained administrative security. Therefore, any resource instance that is not part of a WebSphere Application Server Version 7.0 node cannot be added to an authorization group.

Fine-grained administrative security in a heterogeneous environment has the following restrictions:

- Only nodes that are running WebSphere Application Server Version 7.0 can be part of an administrative authorization group.
- Only servers that are running in a WebSphere Application Server Version 7.0 node can be part of an administrative authorization group.
- Only applications that are targeted on servers running on WebSphere Application Server Version 7.0 can be part of an administrative authorization group.
- If a cluster spans nodes of multiple releases, it cannot be part of an administrative authorization group.

- If a cluster spans nodes of multiple releases, none of its members can be part of an administrative authorization group.
- If an application is targeted on a cluster that spans multiple releases, that application cannot be part of an administrative authorization group.

Fine-grained administrative security in a single-server environment

You can also use fine-grained administrative security in a single-server environment. Various applications in the single server can be grouped and placed in different authorization groups. Therefore, different authorization constraints might exist for different applications.

Life cycle of fine-grained administrative resource

An administrative resource that was once part of an authorization group continues to be part of that authorization group until one of the following events occurs:

- The administrative resource is removed from the authorization group. In this instance, the administrative resource belongs to the cell-level authorization group.
- The administrative resource is removed from the configuration. In this instance, the administrative resource does not exist in the configuration, but still exists in the authorization group. Remove this administrative resource from the authorization group.

After the administrative resource is removed from the authorization group, the administrative authorizer runtime must be notified by using the `AuthorizationManager refreshAll` MBean method.

The **refreshAll** command must be invoked after `AdminConfig.save()` and sync nodes. For example:

JACL:

```
// get AuthorizationGroup Mbean
wsadmin> set agBean [$AdminControl queryNames
Type=AuthorizationGroupManager,process=dmgr,*]
```

JYTHON:

```
// get AuthorizationGroup Mbean
wsadmin> set agBean [$AdminControl queryNames
Type=AuthorizationGroupManager,process=dmgr,*]
```

Example: Using fine-grained security

The following scenarios describe the use of fine-grained administrative security, particularly the new deployment role.

Deployment role scenario 1

In the following scenario, there are four applications configured on server S1, as shown in the following table. Each application must be isolated so that the administrator of one application cannot modify another application. Assume that only user1 can manage application A1, user2 can manage applications A2 and A3, and only user3 can manage application A4.


Note: It is not recommended to have an application in one group and its target server in another group. However, that is not always possible. It is common to have many applications on one server. It is still sometimes necessary to isolate the administration of applications running on the same server.

One example is an Application Service Provider (ASP), where a single application server can have multiple vendor applications. In this instance the server administrator is responsible for installing all

of the vendor applications. Once applications are installed, each vendor can manage their own application without interfering with other vendor's applications.

Application	Server	Node
A1	S1	N1
A2	S1	N1
A3	S1	N1
A4	S1	N1

We can configure authorization groups as shown in the diagram below:



In the diagram, application A1 is in authorization group G1, applications A2 and A3 are in authorization group G2, and application A4 is in authorization group G3.

A deployer role is assigned from authorization group G1 to user1, from authorization group G2 to user2, and from authorization group G3 to user3.


Consequently, user1 can perform all of the operations on application A1, user2 on applications A2 and A3, and user3 on application A4. Since all applications share the same server, we cannot put the same server on all authorization groups. Only a cell-level administrator can install an application. After the installation of an application is complete, the deployer of each application can modify their own. To start and stop the server, cell-level administrative authority is required. This type of scenario is useful in an ASP environment.

Deployment role scenario 2

In the following scenario, a group of applications require the same administrative roles to one server. In this example, applications A1 and A2 are related applications, and can be administrated by one set of administrators. They are running on the same server (S1). Applications A3 and A4 require a different set of administrators, and are running on servers S2 and S3 respectively.

Application	Server	Node
A1	S1	N1
A2	S1	N1

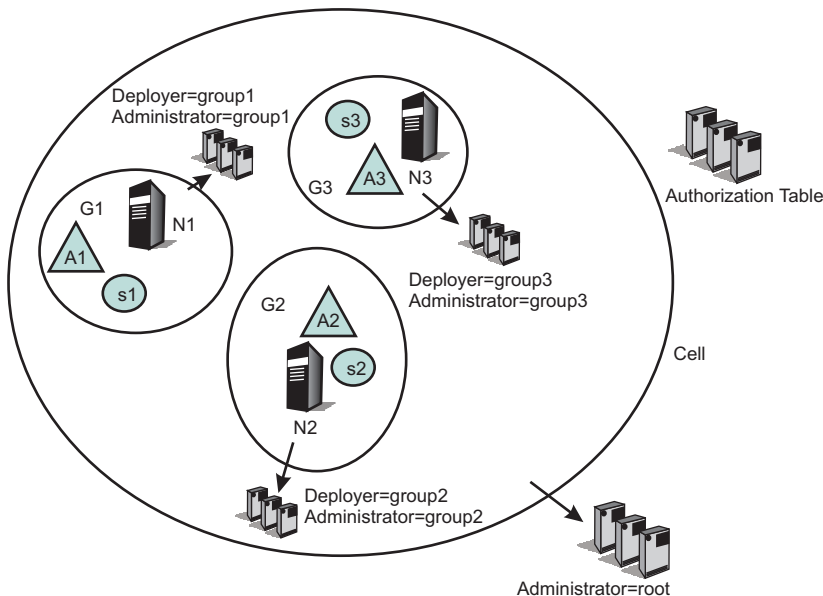
Application	Server	Node
A3	S2	N2
A4	S3	N3



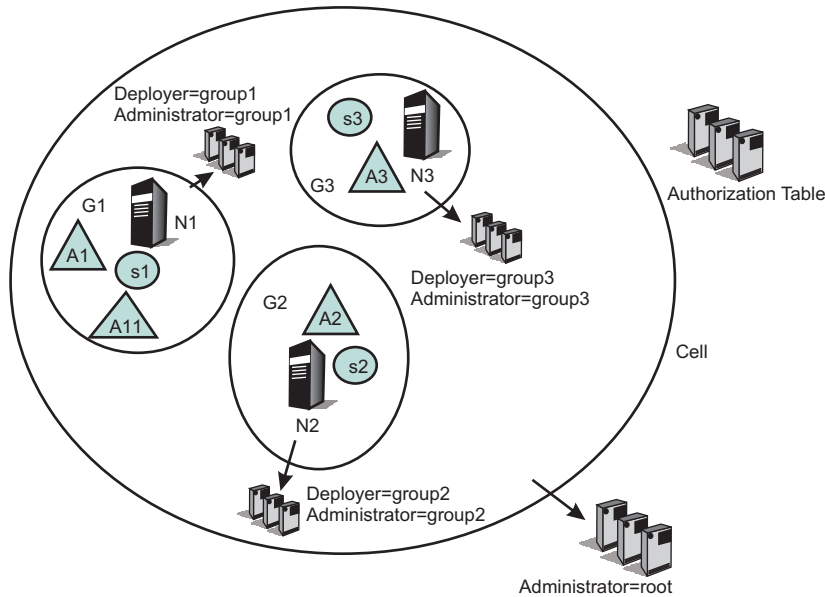
Scenarios that can be applied directly in customer environments

Each developer must be able to modify the configuration for their server, and they must be able to install their application onto that server. They also must be able to start and stop the server as well as the application on the server.

Developers also must be able to configure the server so that they can debug any problems they run into. They must have the ability to update or modify the application being developed. The administrative authorization group for this developer includes at least one server and any applications that the developer installs on that server.



In the following example, developers of authorization group G1 have a new application (A11). They can install and target that new application only on servers within authorization group G1. Also, they can place that new application in their authorization group (G1).



ASP environment scenario


In this scenario, the customer is an ASP. They have their own customers to whom they provide application serving function. They want to enable their customers to administer and monitor their applications, but not to see or administer applications for different customers. In this example, however, the ASP has internal staff administrators whose job it is to maintain the servers.

This internal ASP staff administrator might need to move an application from one server to another to ensure that an application remains available. The internal ASP staff administrator should be able to stop and start the servers and to change their configuration.

In contrast, the ASP customer administrator should not be able to stop or start servers. However, the ASP customer administrator should be able to update their applications running on those servers. The administrative authorization group for the internal ASP administrator can be the whole cell or can include a subset of servers, nodes, clusters and applications. The administrative authorization group for the customer administrator only includes those applications that the customer has paid to have served by this ASP.

When updating the configuration repository, run the admin scripts from the deployment manager so that the fine grain admin security rules will be in effect when admin scripts are run from the deployment manager side.

The following diagram contains a scenario where two different customers have two different types of applications, and can manage their own applications. However, the servers and nodes on which the applications are running are isolated from their customers. The servers and nodes can only be maintained by the internal administrators. In addition, the customers cannot target their applications on a different server. This can only be performed by the internal administrator or internal deployers.




Regional organization scenario


In this scenario, the customer is a large global company. The company's nodes and servers are organized so as to provide application serving for different regions (or alternatively, different lines of business). They want representatives from the different regional areas to be able to monitor and administer the nodes and servers associated with that region. However, they do not want the regional administrator to be able to effect any node and server associated with a different region.

The administrative authorization group for each regional representative includes the nodes, servers, clusters and applications associated with that region.

For example, consider a company that provides multiple services, such as a financial institution that provides services like credit card accounts, brokerage accounts, banking accounts, or travel accounts. Each of these services can be separate applications, and the administrator for each of these applications must also be different. The following figure shows one way to configure such a system:




The following figure shows how the resources in such a system can be grouped to isolate administrators from each other:

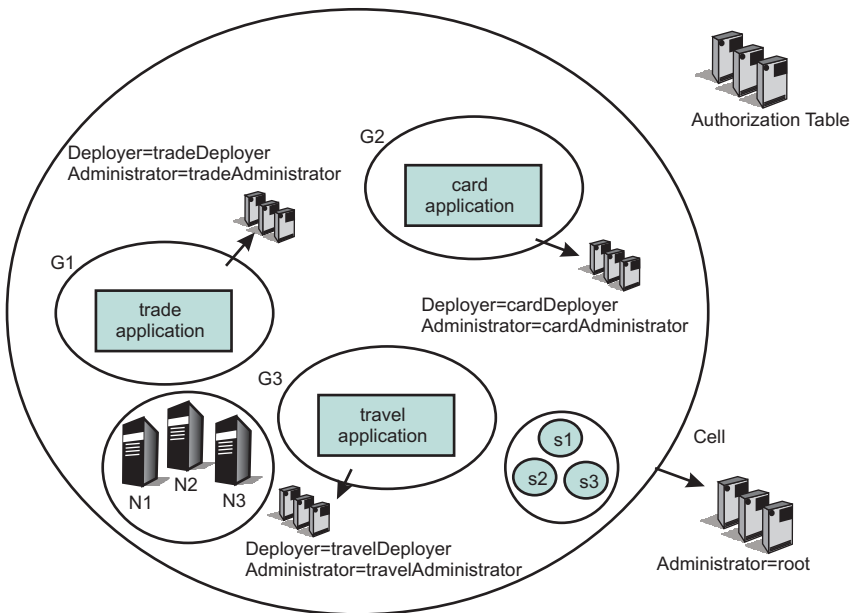


Note that the nodes are not part of any authorization group. Therefore, a trade application administrator cannot stop a server on any of the nodes, and is prevented from stopping a travel application.

The same system can be configured in another way as shown below:



The following figure shows how the resources in such a system can be grouped to isolate administrators from each other:



Chapter 7. Securing communications

WebSphere Application Server provides several methods to secure communication between a server and a client.

About this task

Note: WebSphere Application Server provides several methods for securing communication between a server and a client. New in this release are functions that ensure secure communication between a server and a client. These functions focus on certificate management, authentication, and ensuring trust among the application server, administrative agent, and job manager. The new functions include:

- Creating and using a certificate authority (CA) clients to enable a CA to request, query, and revoke certificates.
- Creating and using chained personal certificates to allow a certificate to be signed with a longer life span.
- Creating and revoking certificate authority (CA) certificates to ensure secure communication between the CA client and the CA server.

The following topics are covered in this section:

- Secure communications using Secure Sockets Layer
- Creating an SSL configuration
- Creating a keystore configuration
- Creating a certificate authority (CA) client
- Deleting a certificate authority (CA) client
- Viewing or Modifying a certificate authority (CA) client
- Creating a keystore configuration for a preexisting keystore file
- Creating a self-signed certificate
- Creating a certificate authority request
- Extracting a signer certificate from a personal certificate
- Retrieving signers from a remote SSL port
- Adding a signer certificate to a keystore
- Adding a signer certificate to the default signers keystore
- Exchanging signer certificates in a keystore
- Configuring certificate expiration monitoring
- Key management for cryptographic uses
- Creating a key set configuration
- Creating a key set group configuration
- Configuring the Web server plug-in for Secure Sockets Layer

Secure communications using Secure Sockets Layer

The Secure Sockets Layer (SSL) protocol provides transport layer security including authenticity, data signing, and data encryption to ensure a secure connection between a client and server that uses WebSphere Application Server. The foundation technology for SSL is *public key cryptography*, which guarantees that when an entity encrypts data using its public key, only entities with the corresponding private key can decrypt that data.

WebSphere Application Server uses Java Secure Sockets Extension (JSSE) as the SSL implementation for secure connections. JSSE is part of the Java 2 Standard Edition (J2SE) specification and is included in the IBM implementation of the Java Runtime Extension (JRE). JSSE handles the handshake negotiation and protection capabilities that are provided by SSL to ensure secure connectivity exists across most protocols. JSSE relies on X.509 certificate-based asymmetric key pairs for secure connection protection and some data encryption. Key pairs effectively encrypt session-based secret keys that encrypt larger blocks of data. The SSL implementation manages the X.509 certificates.

Managing X.509 certificates

Secure communications for WebSphere Application Server require digitally-signed X.509 certificates. The contents of an X.509 certificate, such as its distinguished name and expiration, are either signed by a certificate authority (CA), signed by a root certificate in **NodeDefaultRootStore** or **DmgrDefaultRootStore**, or are self-signed. When a trusted CA signs an X.509 certificate, WebSphere Application Server identifies the certificate and freely distributes it. A certificate must be signed by a CA because the certificate represents the identity of an entity to the general public. Server-side ports that accept connections from the general public must use CA-signed certificates. Most clients or browsers already have the signer certificate that can validate the X.509 certificate so signer exchange is not necessary for a successful connection.

You can trust the identity of a self-signed X.509 certificate only within a peer in a controlled environment, such as internal network communications, accepts the signer certificate. To complete a trusted handshake, you must first send a copy of the entity certificate to every peer that connects to the entity.

CA, chained, and self-signed X.509 certificates reside in Java *keystores*. JSSE provides a reference to the keystore in which a certificate resides. You can select from many types of keystores, including Java Cryptographic Extension (JCE)-based and hardware-based keystores. Typically, each JSSE configuration has two Java keystore references: a keystore and a *truststore*. The keystore reference represents a Java keystore object that holds personal certificates. The truststore reference represents a Java keystore object that holds signer certificates.

A personal certificate without a private key is an X.509 certificate that represents the entity that owns it during a handshake. Personal certificates contain both public and private keys. A signer certificate is an X.509 certificate that represents a peer entity or itself. Signer certificates contain just the public key and verify the signature of the identity that is received during a peer-to-peer handshake.

For more information, see “Extracting a signer certificate from a personal certificate” on page 700

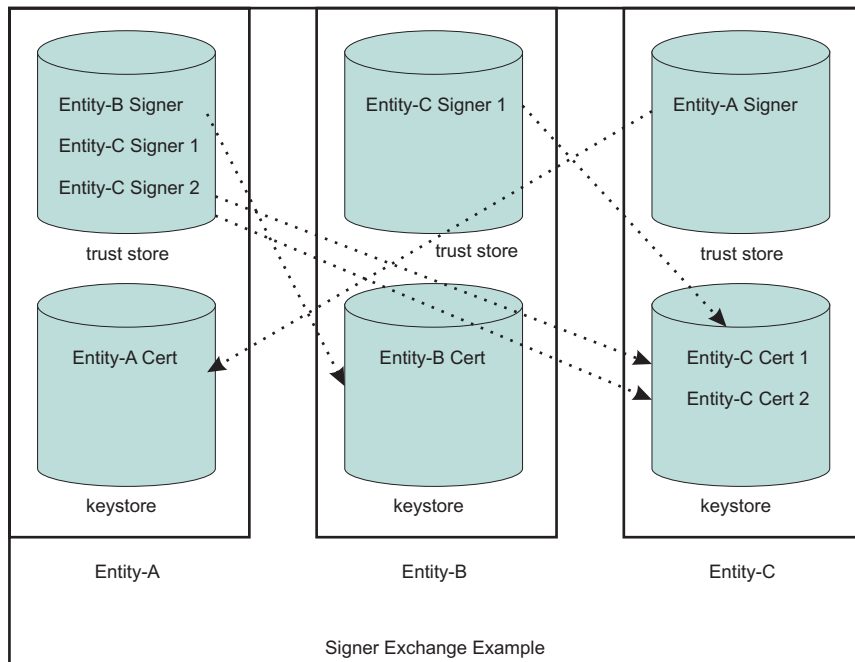
For more information about keystores, see `csec_sslkeystoreconfs.dita`.

Signer exchange

When you configure an SSL connection, you can exchange signers to establish trust in a *personal certificate* for a specific entity. Signer exchange enables you to extract the X.509 certificate from the peer keystore and add it into the truststore of another entity so that the two peer entities can connect. The signer certificate also can originate from a CA as a root signer certificate or a chained certificate’s root signer certificate or an intermediate signer certificate. You can also extract a *signer certificate* directly from a self-signed certificate, which is the X.509 certificate with the public key.

Figure 1 illustrates a hypothetical keystore and truststore configuration. An SSL configuration determines which entities can connect to other entities, and the peer connections that are trusted by an SSL handshake. If you do not have the necessary signer certificate, the handshake fails because the peer cannot be trusted.

Figure 15. Signer exchange



In this example, the truststore for Entity A contains three signers. Entity A can connect to any peer as long as one of the three signers validates its personal certificate. For example, Entity A can connect to Entity B or Entity C because the signers can trust both signed personal certificates. The truststore for Entity-B contains one signer. Entity B is able to connect to Entity C only, and only when the peer endpoint is using certificate Entity-C Cert 1 as its identity. The ports that use the other personal certificate for Entity C are not trusted by Entity B. Entity C can connect to Entity A only.

In the example, the self-signed configuration seems to represent a one-to-one relationship between the signer and the certificate. However, when a CA signs a certificate, it typically signs many at a time. The advantage of using a single CA signer is that it can validate personal certificates that are generated by the CA for use by peers. However, if the signer is a public CA, you must be aware that the signed certificates might have been generated for another company other than your target entity. For your internal communications, private CAs and self-signed certificates are preferable to public CAs because they enable you to isolate the connections that you want to occur and prevent those that you do not want to occur.

SSL configurations

An SSL configuration comprises a set of configuration attributes that you can associate with an *endpoint* or set of endpoints in the WebSphere Application Server topology. The SSL configuration enables you to create an SSLContext object, which is the fundamental JSSE object that the server uses to obtain SSL socket factories. You can manage the following configuration attributes:

- An alias for the SSLContext object
- A handshake protocol version
- A keystore reference
- A truststore reference
- A key manager
- One or more trust managers
- A security level selection of a cipher suite grouping or a specific cipher suite list
- A certificate alias choice for client and server connections

To understand the specifics of each SSL configuration attribute, see “Secure Sockets Layer configurations” on page 595.

Selecting SSL configurations

In previous releases of WebSphere Application Server, you can reference an SSL configuration only by selecting the SSL configuration alias directly. Each secure endpoint was denoted by an alias attribute that references a valid SSL configuration within a repertoire of SSL configurations. When you made a single configuration change, you had to re-configure many alias references across the various processes. Although the current release still supports direct selection, this approach is no longer recommended.

The current release provides improved capabilities for managing SSL configurations and more flexibility when you select SSL configurations. In this release, you can select from the following approaches:

Programmatic selection

You can set an SSL configuration on the running thread prior to an outbound connection. WebSphere Application Server ensures that most system protocols, including Internet Inter-ORB Protocol (IIOP), Java Message Service (JMS), Hyper Text Transfer Protocol (HTTP), and Lightweight Directory Access Protocol (LDAP), accept the configuration. See “Programmatically specifying an outbound SSL configuration using JSSEHelper API” on page 659

Dynamic selection

You can associate an SSL configuration dynamically with a specific target host, port, or outbound protocol by using a predefined selection criteria. When it establishes the connection, WebSphere Application Server checks to see if the target host and port match a predefined criteria that includes the domain portion of the host. Additionally, you can predefine the protocol for a specific outbound SSL configuration and certificate alias selection. See “Dynamic outbound selection of Secure Sockets Layer configurations” on page 607 for more information.

Direct selection

You can select an SSL configuration by using a specific alias, as in past releases. This method of selection is maintained for backwards compatibility because many applications and processes rely on alias references.

Scope selection

You can associate an SSL configuration and its certificate alias, which is located in the keystore associated with that SSL configuration, with a WebSphere Application Server management scope. This approach is recommended to manage SSL configurations centrally. You can manage endpoints more efficiently because they are located in one topology view of the cell. The inheritance relationship between scopes reduces the number of SSL configuration assignments that you must set.

Each time you associate an SSL configuration with a cell scope, the node scope within the cell automatically inherits the configuration properties. However, when you assign an SSL configuration to a node, the node configuration overrides the configuration that the node inherits from the cell. Similarly, all of the application servers for a node automatically inherit the SSL configuration for that node unless you override these assignments. Unless you override a specific configuration, the topology relies on the rules of inheritance from the cell level down to the endpoint level for each application server.

Note: If your applications are relying on SSL configurations that were set as individual settings for each SSL configuration in the topology, but your application servers have inherited the SSL configuration as deployed from the cell level down to the endpoint level, then there is the possibility of communication errors occurring among servers (for example, handshake errors). You need to ensure that your applications are operating consistent with the central management of SSL configurations.

The topology view displays an inbound tree and outbound tree. You can make different SSL configuration selections for each side of the SSL connection based on what that server connects

to as an outbound connection and what the server connects to as an inbound connection. See “Central management of Secure Sockets Layer configurations” on page 608 for more information.

The runtime uses an order of precedence for determining which SSL configuration to choose because you have many ways to select SSL configurations. Consider the following order of precedence when you select a configuration approach:

1. Programmatic selection. If an application sets an SSL configuration on the running thread using the `com.ibm.websphere.ssl.JSSEHelper` application programming interface (API), the SSL configuration is guaranteed the highest precedence.
2. Dynamic selection criteria for outbound host and port or protocol.
3. Direct selection.
4. Scope selection. Scope inheritance guarantees that the endpoint that you select is associated with an SSL configuration and is inherited by every scope beneath it that does not override this selection.

Default chained certificate configuration

By default, WebSphere Application Server creates a unique chained certificate for each node. The chained certificate is signed with a root, a self-signed certificate stored in the **DmgrDefaultRootStore** or **NodeDefaultRootStore**. WebSphere Application Server no longer relies on a self-signed certificate or the default or dummy certificate that is shipped with the product. The `key.p12` default keystore and the `trust.p12` truststore are stored in the configuration repository within the node directory. The default root certificate is stored in the `root-key.p12` in the configuration repository under the node directory.

All of the nodes put their signer certificates from the default root certificate in this common truststore (`trust.p12`). Additionally, after you federate a node, the default SSL configuration is automatically modified to point to the common truststore, which is located in the cell directory. The node can now communicate with all other servers in the cell.

All default SSL configurations contain a keystore with the name suffix `DefaultKeyStore`, a truststore with the name suffix `DefaultTrustStore` and a rootstore with the name suffix `DefaultRootStore`. These default suffixes instruct the WebSphere Application Server runtime to add the root signer of the personal certificate to the common truststore. If a truststore name does not end with `DefaultKeyStore`, the keystores root signer certificates are not added to the common truststore when you federate the server. You can change the default SSL configuration, but you must ensure that the correct trust is established for administrative connections, among others.

For more information, see “Default chained certificate configuration” on page 614 and “Web server plug-in default configuration” on page 730.

Certificate expiration monitoring

Certificate monitoring ensures that the default chained certificate for each node is not allowed to expire. The certificate expiration monitoring function issues a warning before certificates and signers are set to expire. Those certificates and signers that are located in keystores managed by the WebSphere Application Server configuration can be monitored. You can configure the expiration monitor to automatically replace a certificate. A chained certificate will be recreated based on the same data used for the initial creation and sign it with the same root certificate that signed the original certificate. A self-signed certificate or chained certificate is also recreated based upon the same data that is used for the initial creation.

The monitor also can automatically replace old signers with the signers from the new chained or self-signed certificates in keystores that are managed by WebSphere Application Server. The existing signer exchanges that occurred by the runtime during federation and by administration are preserved. For more information, see “Certificate expiration monitoring” on page 622.

The expiration monitor is configured to replace chained personal certificates that are signed by a root certificate in **DmgrDefaultRootStore** or **NodeDefaultRootStore**. The certificate is renewed using the same root certificate that was used to sign the original certificate.

The monitor also can automatically replace old signers with the signers from the new self-signed certificates in keystores that are managed by WebSphere Application Server. The existing signer exchanges that occurred by the runtime during federation and by administration are preserved. For more information, see “Certificate expiration monitoring” on page 622.

WebSphere Application Server clients: signer-exchange requirements

A new chained certificate is generated for each node during its initial startup. To ensure trust, clients must be given the root signers to establish a connection. The introduction of chained certificates in the current release makes this process simpler. Rather than exchanging the signer of a short lived self-signed certificate, you can exchange the long lived root signer which will allow for preserved trust across personal certificate renewals. In addition, you can gain access to the signer certificates of various nodes to which the client must connect with any one of the following options (for more information, see “Secure installation for client signer retrieval” on page 617):

- A signer exchange prompt enables you to import signer certificates that are not yet present in the truststores during a connection to a server. By default, this prompt is enabled for administrative connections and can be enabled for any client SSL configuration. When this prompt is enabled, any connection that is made to a server where the signer is not already present offers the signer of the server along with the certificate information and a Secure Hash Algorithm (SHA) digest of the certificate for verification. The user is given a choice whether to accept these credentials. If the credentials are accepted, the signer is added to the truststore of the client until the signer is explicitly removed. The signer exchange prompt does not occur again when connecting to the same server unless the personal certificate changes.

Note: It is unsafe to trust a signer exchange prompt without verifying the SHA digest. An unverified prompt can originate from a browser when a certificate is not trusted.

- You can run a `retrieveSigners` administrative script from a client prior to making connections to servers. To download signers, no administrative authority is required. To upload signers, you must have Administrator role authority. The script downloads all of the signers from a specified server truststore into the specified client truststore and can be called to download only a specific alias from a truststore. You can also call the script to upload signers to server truststores. When you select the `CellDefaultTrustStore` truststore as the specified server truststore and common truststore for a cell, all of the signers for that cell are downloaded to the specified client truststore, which is typically `ClientDefaultTrustStore`. For more information, see “`retrieveSigners` command” on page 620.
- You can physically distribute to clients the `trust.p12` common truststore that is located in the cell directory of the configuration repository. When doing this distribution, however, you must ensure that the correct password has been specified in the `ssl.client.props` client SSL configuration file. The default password for this truststore is `WebAS`. Change the default password prior to distribution. Physical distribution is not as effective as the previous options. When changes are made to the personal certificates on the server, automated exchange can fail.

Dynamic SSL configuration changes

The SSL runtime for WebSphere Application Server maintains listeners for most SSL connections. A change to the SSL configuration causes the inbound connection listeners to create a new `SSLContext` object. Existing connections continue to use the current `SSLContext` object. Outbound connections automatically use the new configuration properties when they are attempted.

Make dynamic changes to the SSL configuration during off-peak hours to reduce the possibility of timing-related problems and to prevent the possibility of the server starting again. If you enable the runtime

to accept dynamic changes, then change the SSL configuration and save the `security.xml` file. Your changes take effect when the new `security.xml` file reaches each node.

Note: If configuration changes cause SSL handshake failures, administrative connectivity failures also can occur, which can lead to outages. In this case, you must re-configure the SSL connections then perform manual node synchronization to correct the problem. You must carefully complete any dynamic changes. It is highly recommended that you perform changes to SSL configurations on a test environment prior to making the same changes to a production system. For more information, see “Dynamic configuration updates” on page 624.

Built-in certificate management

Certificate management that is comparable to iKeyMan functionality is now integrated into the keystore management panels of the administrative console. Use built-in certificate management to manage personal certificates, certificate requests, and signer certificates that are located in keystores. Additionally, you can remotely manage keystores. For example, you can manage a file-based keystore that is located outside the configuration repository on any node from the deployment manager. You also can remotely manage hardware cryptographic keystores from the deployment manager.

With built-in certificate management, you can replace a chained or self-signed certificate along with all of the signer certificates scattered across many truststores and retrieve a signer from a remote port by connecting to the remote SSL host and port and intercepting the signer during the handshake. The certificate is first validated according to the certificate SHA digest, then the administrator must accept the validated certificate before it can be placed into a truststore.

When you make a certificate request, you can send it to a certificate authority (CA). When the certificate is returned, you can accept it within the administrative console. For more information, see “Certificate management” on page 627.

Note: Although iKeyMan functionality still ships with WebSphere Application Server, configure keystores from the administrative console using the built-in certificate management functionality. iKeyMan is still an option when it is not convenient to use the administrative console. For more information, see `csec_sslkeymancertman.dita`.

AdminTask configuration management

The SSL configuration management panels in the administrative console rely primarily on administrative tasks, which are maintained and enhanced to support the administrative console function. You can use `wsadmin` commands from a Java console prompt to automate the management of keystores, SSL configurations, certificates, and so on.

Secure Sockets Layer configurations

Secure Sockets Layer (SSL) configurations contain attributes that enable you to control the behavior of both the client and the server SSL endpoints. You can assign SSL configurations to have specific management scopes. The scope that an SSL configuration inherits depends upon whether you create it using a cell, node, server, or endpoint link in the configuration topology.

When you create an SSL configuration, you can set the following SSL connection attributes:

- Keystore
- Default client certificate for outbound connections
- Default server certificate for inbound connections
- Truststore
- Key manager for selecting a certificate
- Trust manager or managers for establishing trust during the handshake

- Handshaking protocol
- Ciphers for negotiating the handshake
- Client authentication support and requirements

You can manage an SSL configuration using any of the following methods:

- Central management selection
- Direct reference selection
- Dynamic outbound connection selection
- Programmatic selection

Using the administrative console, you can manage all of the SSL configurations for WebSphere Application Server. From the administrative console, click **Security > SSL certificates and key management > Manage endpoint security configurations > Inbound | Outbound > SSL_configuration**. You can view an SSL configuration at the level it was created and in the inherited scope below that point in the topology. If you want the entire cell to view an SSL configuration, you must create the configuration at the cell level in the topology.

SSL configuration in the security.xml file

The attributes defining an SSL configuration repertoire entry for a specific management scope are stored in the security.xml file. The scope determines the point at which other levels in the cell topology can see the configuration, as shown in the following example:

```
<repertoire xmi:id="SSLConfig_1" alias="NodeDefaultSSLSettings"
managementScope="ManagementScope_1" type="JSSE">
<setting xmi:id="SecureSocketLayer_1" clientAuthentication="false"
clientAuthenticationSupported="false" securityLevel="HIGH" enabledCiphers=""
jsseProvider="IBMJSSE2" sslProtocol="SSL_TLS" keyStore="KeyStore_1"
trustStore="KeyStore_2" trustManager="TrustManager_1" keyManager="KeyManager_1"
clientKeyAlias="default" serverKeyAlias="default"/>
</repertoire>
```

The SSL configuration attributes from the previous code sample are described in Table 1.

Table 21. security.xml Attributes

security.xml attribute	Description	Default	Associated SSL property
xmi:id	The xmi:id attribute represents the unique identifier for this XML entry and determines how the SSL configuration is linked to other XML objects, such as SSLConfigGroup. This system-defined value must be unique.	The administrative configuration service defines the default value.	None. This value is used only for XML associations.
alias	The alias attribute defines the name of the SSL configuration. Direct selection uses the alias attribute and the node is not prefixed to the alias. Rather, the management scope takes care of ensuring that the name is unique within the scope.	The default is CellDefaultSSLSettings.	com.ibm.ssl.alias
managementScope	The managementScope attribute defines the management scope for the SSL configuration and determines the visibility of the SSL configuration at runtime.	The default scope is the cell.	The managementScope attribute is not mapped to an SSL property. However, it confirms whether or not the SSL configuration is associated with a process.

Table 21. security.xml Attributes (continued)

security.xml attribute	Description	Default	Associated SSL property
type	The type attribute defines the Java Secure Socket Extension (JSSE) or System Secure Sockets Layer (SSSL) configuration option. JSSE is the SSL configuration type for most secure communications within WebSphere Application Server.	The default is JSSE.	com.ibm.ssl.sslType
clientAuthentication	The clientAuthentication attribute determines whether SSL client authentication is required.	The default is false.	com.ibm.ssl.clientAuthentication
clientAuthenticationSupported	The clientAuthenticationSupported attribute determines whether SSL client authentication is supported. The client does not have to supply a client certificate if it does not have a client certificate. Note: When you set the clientAuthentication attribute to true, you override the value that is set for the clientAuthenticationSupported attribute.	The default is false.	com.ibm.ssl.client.AuthenticationSupported
securityLevel	The securityLevel attribute determines the cipher suite group. Valid values include HIGH (128-bit ciphers), MEDIUM (40-bit ciphers), LOW (for all ciphers without encryption), and CUSTOM (if the cipher suite group is customized. When you set the enabledCiphers attribute with a specific list of ciphers, the system ignores this attribute.	The default is HIGH.	com.ibm.ssl.securityLevel
enabledCiphers	You can set the enabledCiphers attribute to specify a unique list of cipher suites. Separate each cipher suite in the list with a space.	The default is the securityLevel attribute for cipher suite selection.	com.ibm.ssl.enabledCipherSuites
jsseProvider	The jsseProvider attribute defines a specific JSSE provider.	The default is IBMJSSE2.	com.ibm.ssl.contextProvider
sslProtocol	The sslProtocol attribute defines the SSL handshake protocol. Valid options include: SSLv2 (client-side only), SSLv3, SSL, SSL_TLS, TLSv1, and TLS values. The SSL option includes SSLv2 and SSLv3 values. The TLS option includes the TLSv1 value. SSL_TLS, which is the most interoperable protocol, includes all these values and defaults to a Transport Layer Security (TLS) handshake.	The default is SSL_TLS.	com.ibm.ssl.protocol

Table 21. *security.xml* Attributes (continued)

security.xml attribute	Description	Default	Associated SSL property
keyStore	The keyStore attribute defines the keystore and attributes of the keyStore instance that the SSL configuration uses for key selection.	The default is CellDefaultKeyStore.	For more information, see Keystore configurations.
trustStore	The trustStore attribute defines the key store that the SSL configuration uses for certificate signing verification.	The default is CellDefaultTrustStore.	A trustStore is a logical JSSE term. It signifies a key store that contains signer certificates. Signer certificates validate certificates that are sent to WebSphere Application Server during an SSL handshake.
keyManager	The keyManager attribute defines the key manager that WebSphere Application Server uses to select keys from a key store. A JSSE key manager controls the javax.net.ssl.X509KeyManager interface. A custom key manager controls the javax.net.ssl.X509KeyManager and the com.ibm.wsspi.ssl.KeyManagerExtendedInfo interfaces. The com.ibm.wsspi.ssl.KeyManagerExtendedInfo interface provides more information from WebSphere Application Server.	The default is IbmX509.	com.ibm.ssl.keyManager defines a well-known key manager and accepts the algorithm and algorithm/provider formats, for example IbmX509 and IbmX509 IBMJSSE2. com.ibm.ssl.customKeyManager defines a custom key manager and takes precedence over the other keyManager properties. This class must implement javax.net.ssl.X509KeyManager and can implement com.ibm.wsspi.ssl.KeyManagerExtendedInfo. For more information, see csec_sslx509certIDkeyman.dita
trustManager	The trustManager determines which trust manager or list of trust managers to use for determining whether to trust the peer side of the connection. A JSSE trust manager implements the javax.net.ssl.X509TrustManager interface. A custom trust manager might also implement com.ibm.wsspi.ssl.TrustManagerExtendedInfo interface to get more information from the WebSphere Application Server environment.	The default is IbmX509. You can specify the IbmPKIX trust manager for certificate revocation list (CRL) verification when the certificate contains a CRL distribution point.	com.ibm.ssl.trustManager defines a well-known trust manager, which is required for most handshake situations. com.ibm.ssl.trustManager performs certificate expiration checking and signature validation. You can define com.ibm.ssl.customTrustManagers with additional custom trust managers that are called during an SSL handshake. Separate additional trust managers with the vertical bar () character. For more information, see csec_sslx509certtrustdecisions.dita

Client SSL configurations are managed using the `ssl.client.props` properties file. The `ssl.client.props` file is located in the `${USER_INSTALL_ROOT}/properties` directory for each profile. For more information about configuring this file, see the “`ssl.client.props` client configuration file” on page 670.

Trust manager control of X.509 certificate trust decisions

The role of the trust manager is to validate the Secure Sockets Layer (SSL) certificate that is sent by the peer, which includes verifying the signature and checking the expiration date of the certificate. A Java Secure Socket Extension (JSSE) trust manager determines if the remote peer can be trusted during an SSL handshake.

WebSphere Application Server has the ability to call multiple trust managers during an SSL connection. The default trust manager does the standard certificate validation; custom trust manager plug-ins run customized validation such as host name verification. For more information, see “Example: Developing a custom trust manager for custom SSL trust decisions” on page 650

When a trust manager is configured in a server-side SSL configuration, the server calls the `isClientTrusted` method. When a trust manager is configured in a client-side SSL configuration, the client calls the

isServerTrusted method. The peer certificate chain is passed to these methods. If the trust manager chooses not to trust the peer information, it might produce an exception to force a handshake failure.

Optionally, WebSphere Application Server provides the `com.ibm.wsspi.ssl.TrustManagerExtendedInfo` interface so that additional information can be passed to the trust manager. For more information, see the `com.ibm.wsspi.ssl.TrustManagerExtendedInfo` interface.

Default IbmX509 trust manager

The default IbmX509 trust manager, which is used in the following code sample, establishes trust by performing standard certificate validation.

```
<trustManagers xmi:id="TrustManager_1132357815717" name="IbmX509" provider="IBMJSSE2"
algorithm="IbmX509" managementScope="ManagementScope_1132357815717"/>
```

The trust manager provides a signer certificate to verify the peer certificate that is sent during the handshake. The signers who are added to the truststore for the SSL configuration must be trustworthy. If you do not trust the signers or do not want to allow others to connect to your servers, consider removing default root certificates from certificate authorities (CA). You might also remove any certificates if you cannot verify their origination.

Default IbmPKIX trust manager

You can use the default IbmPKIX trust manager to replace the IbmX509 trust manager, which is shown in the following code sample:

```
<trustManagers xmi:id="TrustManager_1132357815717" name="IbmPKIX" provider="IBMJSSE2"
algorithm="IbmPKIX" trustManagerClass="" managementScope="ManagementScope_1132357815717">
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_1132357815717"
name="com.ibm.security.enableCRLDP" value="true" type="boolean"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_1132357815718"
name="com.ibm.jsse2.checkRevocation" value="true" type="boolean"/>
</trustManagers>
```

```
<trustManagers xmi:id="TrustManager_managementNode_2" name="IbmPKIX" provider=
"IBMJSSE2" algorithm="IbmPKIX" trustManagerClass=""
managementScope="ManagementScope_managementNode_1">
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_1" name="com.ibm.se
curity.enableCRLDP" value="false" type="boolean" displayNameKey="" nlsRangeKey=""
hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_2" name="com.ibm.js
se2.checkRevocation" value="false" type="boolean" displayNameKey="" nlsRangeKey=
"" hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_3" name="ocsp.enabl
e" value="false" type="String" displayNameKey="" nlsRangeKey="" hoverHelpKey=""
range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_4" name="ocsp.respo
nderURL" value="http://ocsp.example.net:80" type="String" displayNameKey=""
nlsRangeKey="" hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_5" name="ocsp.respo
nderCertSubjectName" value="" type="String" displayNameKey="" nlsRangeKey="" hov
erHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_6" name="ocsp.respo
nderCertIssuerName" value="" type="String" displayNameKey="" nlsRangeKey="" hove
rHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_7" name="ocsp.respo
nderCertSerialNumber" value="" type="String" displayNameKey="" nlsRangeKey="" ho
verHelpKey="" range="" inclusive="false" firstClass="false"/>
</trustManagers>
```

See “Example: Enabling certificate revocation checking with the default IbmPKIX trust manager” on page 601 for additional information in using the default IbmPKIX trust manager.

In addition to its role of standard certificate verification, the IbmPKIX trust manager checks for OCSP properties and for certificates that contain certificate revocation list (CRL) distribution points. This process is known as extended CRL checking. When you select a trust manager, its associated properties are automatically set as Java System properties so that the IBM CertPath and IBMJSSE2 providers are aware that CRL checking is enabled.

Custom trust manager

You can define a custom trust manager to perform additional trust checking, which is based upon the needs of the environment. For example, in one environment, you might enable connections from the same Transmission Control Protocol (TCP) subnet only. The `com.ibm.wsspi.ssl.TrustManagerExtendedInfo` interface provides extended information about the connection that is not provided by the standard Java Secure Sockets Extension (JSSE) `javax.net.ssl.X509TrustManager` interface. The configured `trustManagerClass` attribute determines which class is instantiated by the runtime, as shown in the following code sample:

```
<trustManagers xmi:id="TrustManager_1132357815718" name="CustomTrustManager"
trustManagerClass="com.ibm.ws.ssl.core.CustomTrustManager"
managementScope="ManagementScope_1132357815717"/>
```

The `trustManagerClass` attribute must implement the `javax.net.ssl.X509TrustManager` interface and, optionally, can implement the `com.ibm.wsspi.ssl.TrustManagerExtendedInfo` interface.

Disabling the default trust manager

In some cases, you might not want to perform the standard certificate verification that is provided by the IbmX509 and IbmPKIX default trust managers. For example, you might be working with an internal automated test infrastructure that is not concerned with SSL client or server authentication, integrity, or confidentiality. The following sample code shows a basic custom trust manager such as `com.ibm.ws.ssl.core.CustomTrustManager` whose property is set to `true`.

```
com.ibm.ssl.skipDefaultTrustManagerWhenCustomDefined=true
```

You can set this property in the global properties at the top of the `ssl.client.props` file for clients or in the `security.xml` custom properties file for servers. You must configure a custom trust manager when you disable the default trust manager to prevent the server from calling the default trust manager even though it is configured. Disabling the default trust manager is not a common practice. Be sure to test the system with the disabled default trust manager in a test environment first. For more information on setting up a custom trust manager, see “Creating a custom trust manager configuration” on page 646

Key manager control of X.509 certificate identities

The role of a Java Secure Socket Extension (JSSE) key manager is to retrieve the certificate that is used to identify the client or server during a Secure Sockets Layer (SSL) handshake.

WebSphere Application Server provides a default key manager that can select a certificate from a keystore when you define the following SSL configuration properties:

com.ibm.ssl.keyStoreClientAlias

Defines the alias that is chosen from the keystore for the client side of a connection. This alias must be present in the keystore.

com.ibm.ssl.keyStoreServerAlias

Defines the alias that is chosen from the keystore for the server side of a connection. This alias must be present in the keystore.

These two properties are set automatically when you use the administrative console because the default key manager is already configured.

With WebSphere Application Server, you can configure only one key manager at a time for a given SSL configuration. If you want custom certificate selection logic on the client side, you must write a new custom key manager. The custom key manager could provide function that prompts the user to choose a certificate dynamically. Also, you can implement an extended interface so that a key manager can provide information during connection time. For more information on the extended interface, see the `com.ibm.wsspi.ssl.KeyManagerExtendedInfo` interface. For more information on custom key manager development, see `rsec_ssldevcustomkeymgr.dita`.

Default IbmX509 key manager

The default IbmX509 key manager chooses a certificate to serve as the identity for an SSL handshake. The key manager is called to enable client authentication on either side of the SSL handshake; frequently on the server-side, and less frequently on the client side according to client and server requirements. If a keystore is not configured on the client-side and SSL client authentication is enabled, the key manager cannot select a certificate to send to the server. Therefore, the handshake fails.

The following sample code shows the key manager configuration in the `security.xml` file for an IbmX509 key manager.

```
<keyManagers xmi:id="KeyManager_1" name="IbmX509"
provider="IBMJSSE2" algorithm="IbmX509" keyManagerClass=""
managementScope="ManagementScope_1"/>
```

You do not specify the `keyManagerClass` class because the key manager is provided by the IBMJSSE2 provider. However, you can specify whether the key manager is a custom class implementation, in which case you must specify the `keyManager` class, or an algorithm name that WebSphere Application Server can start from the Java security provider framework.

Custom key manager

The following sample code shows the key manager configuration in the `security.xml` file for a custom class.

```
<keyManagers xmi:id="KeyManager_2" name="CustomKeyManager"
keyManagerClass="com.ibm.ws.ssl.core.CustomKeyManager"
managementScope="ManagementScope_1"/>
```

The custom class must implement the `javax.net.ssl.X509KeyManager` interface and, optionally, implement the `com.ibm.wsspi.ssl.KeyManagerExtendedInfo` interface to retrieve additional WebSphere Application Server information. This interface replaces the function of the default key manager because you can configure only one key manager at a time. Therefore, the custom key manager has sole responsibility for selecting the alias to use from the configured keystore. The benefit of a custom key manager is its ability, on the client side, to prompt for an alias. This process enables the user to decide which certificate to use in situations where the user knows the client certificate identity. For more information, see `tsec_sslcreatecuskeymgr.dita`.

Example: Enabling certificate revocation checking with the default IbmPKIX trust manager

The IbmPKIX trust manager is enabled in the WebSphere Application Server by default. The IbmPKIX trust manager allows certificate revocation checking to occur. You enable certificate revocation checking by using the administrative console or by manually updating the `ssl.client.props` file.

The default IbmPKIX trust manager

The IbmPKIX trust manager is enabled by default, but revocation checking is not enabled by default. The following trust manager definition for IbmPKIX reflects the default condition:

```
<trustManagers xmi:id="TrustManager_managementNode_2" name="IbmPKIX" provider=
"IBMJSSE2" algorithm="IbmPKIX" trustManagerClass=""
managementScope="ManagementScope_managementNode_1">
```

```

<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_1" name="com.ibm.se
curity.enableCRLDP" value="false" type="boolean" displayNameKey="" nlsRangeKey=""
" hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_2" name="com.ibm.js
se2.checkRevocation" value="false" type="boolean" displayNameKey="" nlsRangeKey=""
" hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_3" name="ocsp.enabl
e" value="false" type="String" displayNameKey="" nlsRangeKey="" hoverHelpKey=""
range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_4" name="ocsp.respo
nderURL" value="http://ocsp.example.net:80" type="String" displayNameKey=""
nlsRangeKey="" hoverHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_5" name="ocsp.respo
nderCertSubjectName" value="" type="String" displayNameKey="" nlsRangeKey="" hov
erHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_6" name="ocsp.respo
nderCertIssuerName" value="" type="String" displayNameKey="" nlsRangeKey="" hove
rHelpKey="" range="" inclusive="false" firstClass="false"/>
<additionalTrustManagerAttrs xmi:id="DescriptiveProperty_7" name="ocsp.respo
nderCertSerialNumber" value="" type="String" displayNameKey="" nlsRangeKey="" ho
verHelpKey="" range="" inclusive="false" firstClass="false"/>
</trustManagers>

```

Enabling certificate revocation checking with the default IbmPKIX trust manager

You can view and change IbmPKIX Trust Manager Custom Properties using the administrative console.

To do this,

- Click **Security > SSL certificate and key management**.
- Under Related Items, click **Trust managers**.
- Click **IbmPKIX**.
- Under Additional Properties, click **Custom properties**.

IbmPKIX custom properties

com.ibm.jsse2.checkRevocation

This property configures revocation checking for the Java Virtual Machine (JVM). This property is set to false by default because the default WebSphere certificates used for SSL communication do not contain certificate revocation list (CRL) distribution points or Online Certificate Status Protocol (OCSP) information.

default

false

com.ibm.security.enableCRLDP

This property configures CRL distribution point checking for the PKIX trust manager.

Note: If you enable CRL distribution point revocation checking, the certificates used for secure sockets layer (SSL) must contain a valid distribution point and the distribution point must be accessible or else SSL communication will fail and the server will not function correctly.

default

false

For certificates that do not contain an internal CRL distribution point, the following properties can be used so the revocation status will be checked against a remote LDAP server containing the CRL.

com.ibm.security ldap.certstore.host

This property specifies the LDAP server host name containing trusted certificates or certificate revocation lists. The target LDAP server host is used to obtain CA certificates or certificate revocation lists when validating a certificate and the local truststore does not contain the required

certificate. The local truststore must contain the required certificates if an LDAP server is not specified. In cases when an LDAP server is used, the root CA certificates must also be located in the local truststore as the LDAP server is not a trusted certificate store.

Note: Enabling this property in addition to the `com.ibm.jsse2.checkRevocation` property enables revocation checking. The remote LDAP server must contain a valid certificate revocation list and the server must be accessible. If the revocation status cannot be determined then the check will fail and SSL communication will fail and the server will not function correctly.

default
none

com.ibm.security.ldap.certstore.port

This property specifies the LDAP server port. A port value of 389 will be used by default if no LDAP server port is specified.

default
389

The following Java Development Kit (JDK) properties apply to enabling certificate revocation checking with the default `IbmPKIX` trust manager:

- `ocsp.enable`
- `ocsp.responder`
- `ocsp.responderCertSubjectName`
- `ocsp.responderCertIssuerName`
- `ocsp.responderCertSerialNumber`

These JDK properties can be set using the administrative console. You should reference *Java(TM) Certification Path API Programmer's Guide - SDK 6.0* for descriptions of these properties and their allowable settings.

Note: In addition to its role of standard certificate verification, the `IbmPKIX` trust manager checks for certificates that contain CRL distribution points. This process is known as extended CRL checking. By default, CRL distribution point revocation checking is disabled. To enable CRL distribution point revocation checking, you must set the following properties to `true` using the administrative console:

- `com.ibm.security.enableCRLDP`
- `com.ibm.jsse2.checkRevocation`

OCSP properties and CRL properties affect certificate revocation checking. By default OCSP properties are checked first. If there is an error validating the certificate with OCSP, then validation uses a CRL distribution point instead.

When you select a trust manager, its associated properties are automatically set as Java system properties so that the `IBMCertPath` and `IBMJSSE2` providers are aware that CRL checking is enabled or disabled. Similarly, the same applies for OCSP properties, which are `java.security.Security` properties.

Client considerations

You can also enable revocation checking for WebSphere application and administrative clients by directly setting the properties in the `ssl.client.props` file. An example of the `ssl.client.props` file follows:

```
#-----  
# Default Revocation Checking Properties  
# These properties are used for certificate revocation checking with the IBM  
# PKIX TrustManager.  
#
```

```

# To enable CRL Distribution Points extension checking, use the system property
# com.ibm.security.enableCRLDP.
#
# OCSP checking is not enabled by default. It is enabled by setting the
# ojsp.enable property to "true". Use of the other ojsp properties is optional.
#
# Note: Both OCSP and CRLDP checking is only effective if revocation checking
# has also been enabled by setting com.ibm.jsse2.checkRevocation to "true".
#
#-----
com.ibm.jsse2.checkRevocation=false
com.ibm.security.enableCRLDP=false
#ocsp.enable=true
#ocsp.responderURL=http://ocsp.example.net
#ocsp.responderCertSubjectName=CN=OCSP Responder, O=XYZ Corp
#ocsp.responderCertIssuerName=CN=Enterprise CA, O=XYZ Corp
#ocsp.responderCertSerialNumber=2A:FF:00

```

Note: In order for these properties to be effective, you must ensure that the IbmPKIX trust manager is initialized by setting `com.ibm.ssl.trustManager=IbmPKIX`.

In addition, for revocation checking to be processed successfully on the client, you are required to turn off the signer exchange prompt. To do this, change the value of the **com.ibm.ssl.enableSignerExchangePrompt** property to `false`, in the `ssl.client.props` file.

For more information on these properties, see *Java(TM) Certification Path API Programmer's Guide - SDK 6.0*.

Keystore configurations

Use keystore configurations to define how the runtime for WebSphere Application Server loads and manages keystore types for Secure Sockets Layer (SSL) configurations.

By default, the `java.security.Security.getAlgorithms("KeyStore")` attribute does not display a predefined list of keystore types in the administrative console. Instead, WebSphere Application Server retrieves all of the KeyStore types that can be referenced by the `java.security.KeyStore` object, including hardware cryptographic, z/OS platform, i5/OS platform, IBM Java Cryptography Extension (IBMJCE), and Java-based content management system (CMS)-provider keystores. If you specify a keystore provider in the `java.security` file or add it to the provider list programmatically, WebSphere Application Server also retrieves custom keystores. The retrieval list depends upon the `java.security` configuration for that platform and process.

IBMJCE file-based keystores (JCEKS, JKS, and PKCS12)

A typical IBMJCE file-based keystore configuration is shown in the following sample code:

```

<keyStores xmi:id="KeyStore_1" name="NodeDefaultKeyStore"
password="{xor}349dkckdd=" provider="IBMJCE"
location="{USER_INSTALL_ROOT}/config/cells/myhostNode01Cell
/nodes/myhostNode01/key.p12" type="PKCS12" fileBased="true"
hostList="" initializeAtStartup="true" readOnly="false"
description="Default key store for myhostNode01" usage="SSLKeys"
managementScope="ManagementScope_1"/>

```

For more information about default keystore configurations, see “Default chained certificate configuration” on page 614.

Table 1 describes the attributes that are used in the sample code.

Table 22. keystore configurations

Attribute name	Default	Description
xmi:id	Varies	A value that issued to reference the keystore from another area in the configuration, for example, from an SSL configuration. Make this value unique within the security.xml file.
name	For JSSE keystore: CellDefaultKeyStore. For JSSE truststore: CellDefaultTrustStore.	A name that is used to identify the keystore by sight. The name can determine if the keystore is a default keystore based upon whether the name ends with DefaultKeyStore or DefaultTrustStore.
password	The default keystore password is WebAS. It is recommended that this be changed as soon as possible. See "Updating default key store passwords using scripting" on page 963 for more information.	The password that is used to access the keystore name is also the default that is used to store keys within the keystore.
description	No default	A description of the keystore.
usage	An attribute specifying what the keystore is used for.	Valid values are: SSLKeys, KeySetKeys, RootKeys, DeletedKeys, DefaultSigners, RSATokenKeys.
provider	The default provider is IBMJCE.	The Java provider that implements the type attribute (for example, PKCS12 type). The provider can be left unspecified and the first provider that implements the keystore type specified is used.
location	The default varies, but typically references a key.p12 file or a trust.p12 file in the node or cell directories of the configuration repository. These files are PKCS12 type keystores.	The keystore location reference. If the keystore is file-based, the location can reference any path in the file system of the node where the keystore is located. However, if the location is outside of the configuration repository, and you want to manage the keystore remotely from the administrative console or from the wsadmin utility, then specify the hostList attribute that contains the host name of the node where it resides.
type	The default Java crypto device keystore type is PKCS12.	This type specifies the keystore. Valid types can be those returned by the <code>java.security.Security.getAlgorithms("KeyStore")</code> attribute. These types include the following keystore types, and availability depends on the process and platform <code>java.security</code> configuration: <ul style="list-style-type: none"> • JKS • JCEKS • PKCS12 • PKCS11 (Java crypto device) • CMSKS • IBMi5OSKeyStore • JCERACFKS • JCECAKS keystores (replacing JCE4758KS) - (z/OS crypto device)
fileBased	The default is true.	This option is required for default keystores. It indicates a file-system keystore so you can use a <code>FileInputStream</code> or <code>FileOutputStream</code> for loading and storing the keystore.

Table 22. keystore configurations (continued)

Attribute name	Default	Description
hostList	The hostList attribute is used to specify a remote hostname so that the keystore can be remotely managed. There are no remotely managed keystores by default. All default keystores are managed locally in the configuration repository and synchronized out to each of the nodes.	The option manages a keystore remotely. You can set the host name of a valid node for a keystore. When you use either the administrative console or the wsadmin utility to manage certificates for this keystore, an MBean call is made to the node where the keystore exists for the approved operation. You can specify multiple hosts, although synchronization of the keystore operations are not guaranteed. For example, one of the hosts that is listed might be down when a specific operation is performed. Therefore, use multiple hosts in this list.
initializeAtStartup	The default is true.	This option informs the runtime to initialize the keystore during startup. This option can be important for hardware cryptographic device acceleration.
readOnly	The default is false.	This option informs the configuration that you cannot write to this keystore. That is, certain update operations on the keystore cannot be attempted and are not allowed. An example of a read-only keystore type is JCERACFKS on the z/OS platform. This type is read-only from the WebSphere certificate management standpoint, but you can also update it using the keystore management facility for RACF.
managementScope	The default scope is the node scope for a base Application Server environment and the cell scope for a Network Deployment environment.	This option references a particular management scope in which you can see this keystore. For example, if a hardware cryptographic device is physically located on a specific node, then create the keystore from a link to that node in the topology view under Security > Security Communications > SSL configurations . You can also use management scope to isolate a keystore reference. In some cases, you might need to allow only a specific application server to reference the keystore; the management scope is for that specific server.

CMS keystores

You can set some provider-specific attributes in CMS keystores.

If the CMSKS provider supports the createStashFileForCMS attribute, and you set the attribute to true for CMSKS keystores, WebSphere Application Server creates an .sth file in the keystore location that is referenced by the attribute. The .sth extension is appended to the keystore name. For example, if the CMSKS keystore is available for a plug-in configuration and you set createStashFileForCMS to true, the stash file that is represented in the following sample code is created in the \${USER_INSTALL_ROOT}\profiles\AppSrv01/config/cells/myhostCell01/nodes/myhostNode01/servers/webserver1/plugin-key.sth path.

```
<keyStores xmi:id="KeyStore_1132071489571" name="CMSKeyStore"
password="{xor}HRYNFATrbxEw0zpvbhw6MzM=" provider="IBMCMSProvider"
location="${USER_INSTALL_ROOT}\profiles\AppSrv01/config/cells/myhostCell01
/nodes/myhostNode01/servers/webserver1/plugin-key.kdb" type="CMSKS"
fileBased="true" createStashFileForCMS="true"
managementScope="ManagementScope_1132071489569"/>
```

When you create a CMS keystore, the CMS provider is IBMi50SJSSEProvider, and the CMS type is IBMi50SKeyStore, as shown in the following sample code:


```
<keyStores xmi:id="KeyStore_1132071489571" name="CMSKeyStore"
password="{xor}HRYNFAtRbxEw0zpbhw6MzM=" provider="IBMi50JSSEProvider"
location="{USER_INSTALL_ROOT}\profiles\AppSrv01/config/cells/myhostCell01
/nodes/myhostNode01/servers/webserver1/plugin-key.kdb" type="IBMi50SKeyStore"
fileBased="true" createStashFileForCMS="true"
managementScope="ManagementScope_1132071489569"/>
```

Hardware cryptographic keystores

For cryptographic device configuration, see “Key management for cryptographic uses” on page 716 and “Configuring a hardware cryptographic keystore” on page 680.

You can add a slot either as the custom property, `com.ibm.ssl.keyStoreSlot`, or as the configuration attribute, `slot="0"`. The custom property is read before the attribute for backwards compatibility.

In certain environments, you can use the hardware cryptographic card for hardware acceleration. Either set the `useForAcceleration` attribute to `true` or set the `com.ibm.ssl.keyStoreUseForAcceleration` custom property. When you set the attribute to `true`, you are not required to configure a password. However, you cannot use the device to store keys. For more information on configuring a hardware cryptographic keystore, see [Configuring hardware cryptographic keystore configurations](#).

Dynamic outbound selection of Secure Sockets Layer configurations

WebSphere Application Server provides dynamic outbound selection that enables you to choose a specific Secure Sockets Layer (SSL) configuration and certificate alias for each outbound protocol, target host, target port, or any combination of these attributes. You can specify the dynamic selection information for outbound connections from a pure client or from a server that is acting as a client.

Before the SSL runtime for WebSphere Application Server starts an outbound connection, the runtime attempts to match the outbound protocol, target host, and target port attributes with the dynamic outbound selection information that is associated with an SSL configuration and certificate alias in the configuration.

The runtime caches both selection misses and selection hits, so the impact on performance can be minimal. However, a relationship exists between the amount of dynamic outbound selection information and its impact on the initial connection performance.

Dynamic selection information for outbound administrative SOAP connections

The default dynamic outbound selection configuration specifies that all outbound administrative SOAP connections use the same SSL configuration and certificate alias within the entire cell.

The management scope of the SSL configuration is at the cell level so every process in the cell sees this configuration. When an administrative SOAP client makes an outbound connection, the client provides protocol, target host, and target port information to the SSL runtime, as shown in the following sample code:

```
<dynamicSSLConfigSelections xmi:id="DynamicSSLConfigSelection_1"
name="AdminSSLOutbound" description="Uses the AdminSSL configuration
for all outbound SOAP calls." dynamicSelectionInfo="ADMIN_SOAP,*,*"
certificateAlias="default" sslConfig="SSLConfig_2"
managementScope="ManagementScope_1"/>
```

The SSL runtime checks the configuration and confirms that `certificateAlias="default"` and `sslConfig="SSLConfig_2"` for the `ADMIN_SOAP` protocol when you enable client certificate authentication.

Note: The configured target host and port are represented by the `*` character because both the host and the port attributes can be any value.

Target information during outbound connections

The dynamic outbound selection configurations are only effective when the outbound protocol uses the JSSEHelper application programming interface (API) when you select an SSL configuration with a specified connectionInfo hash map. This hash map must contain the following properties:

com.ibm.ssl.direction

The value for outbound connections is OUTBOUND.

com.ibm.ssl.remoteHost

The format should match what the protocol provides. Typically this is the canonical Domain Name Space (DNS), but it also could be the IP address.

com.ibm.ssl.remotePort

The port is target port.

com.ibm.ssl.endPointName

The value for an outbound connection must be one of the following protocol strings:

- IIOP
- HTTP
- SIP
- LDAP
- ADMIN_IPC
- ADMIN_SOAP
- BUS_TO_BUS
- BUS_CLIENT
- BUS_TO_WEBSPHERE_MQ

Central management of Secure Sockets Layer configurations

By default, Secure Sockets Layer (SSL) configurations for servers are managed from a central location in the topology view of the administrative console. You can associate an SSL configuration and certificate alias with a specific management scope. This method is the most efficient method to manipulate and modify configurations when the server topology changes.

In prior releases, SSL configurations are managed for each process. You have to maintain individual settings for each SSL configuration in the topology. In this release of WebSphere Application Server, management control of your SSL configurations offers more options and additional flexibility. You are able to make coarse-grained changes for the entire topology using the cell-scope and also make fine-grained changes using a particular endpoint name for a specific application server process. Because the SSL configuration associations manifest an inheritance behavior, you can simplify the number of associations by referencing only the highest level management scope that needs a unique configuration.

The topology view provides the scoping mechanism. The SSL configuration inherits its scope, which can be seen as its display in the topology. The scope encompasses the level where you created the configuration and all the levels below that point. For example, when you create an SSL configuration at a specific node, that configuration can be seen by that node agent and by every application server that is part of that node. Any application server or node that is not part of this particular node can not see this SSL configuration.

Your security environment influences issues such as the uniqueness of the SSL configurations, as well as the SSL configuration and the certificate alias placement in the topology. You are also able to configure different certificate aliases and different SSL configurations for inbound connections versus outbound connections.

To configure the inbound and outbound topologies, which must be done separately in the administrative console, click **Security > SSL certificates and key management > Manage endpoint security configurations > Inbound | Outbound**.

Default centrally managed SSL configuration

It is simpler to manage SSL configurations centrally in the topology view of the administrative console, but you can also use wsadmin scripting within AdminTasks to manage SSL configurations.

The configuration element of the `security.xml` file can be used to manage SSL configuration associations. The `sslConfigGroup` configuration object is the mechanism that is used to associate a connection direction and management scope with a specific SSL configuration and certificate alias. The default `sslConfigGroups` cell attribute has a predefined inbound and outbound cell-scoped configuration that each endpoint in the cell inherits. Because rules of precedence must guide you when you select SSL configurations, see `csec_sslsecurecom.dita` before you modify your configurations using wsadmin scripting.

```
<sslConfigGroups xmi:id="SSLConfigGroup_1"
name="myhostCell01" direction="inbound" certificateAlias="default"
sslConfig="SSLConfig_1" managementScope="ManagementScope_1"/>
<sslConfigGroups xmi:id="SSLConfigGroup_2" name="myhostCell01"
direction="outbound" certificateAlias="default" sslConfig="SSLConfig_1"
managementScope="ManagementScope_1"/>

<managementScopes xmi:id="ManagementScope_1"
scopeName="(cell):myhostCell01" scopeType="cell"/>
```

In the previous sample code, the `sslConfigGroups` attribute references the cell management scope. For that example, if a different scope was intended the following list shows the precedence order for selecting management scopes, from the highest order of precedence to the lowest. Each time an endpoint scope is defined, it uses the specified SSL configuration and certificate alias.

Endpoint scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01:
(node):myhostNode01:(server):server1:(endpoint):ENDPOINT_NAME_IN_SERVERINDEX"
scopeType="endpoint"/>
```

Server scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01:
(node):myhostNode01:(server):server1" scopeType="server"/>
```

Cluster scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01:
(cluster):myCluster" scopeType="cluster"/>
```

Node scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01:
(node):myhostNode01" scopeType="node"/>
```

Nodegroup scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01:
(nodegroup):DefaultNodeGroup" scopeType="nodegroup"/>
```

Cell scope

```
<managementScopes xmi:id="ManagementScope_1" scopeName="(cell):myhostCell01"
scopeType="cell"/>
```

Secure Sockets Layer node, application server, and cluster isolation

Secure Sockets Layer (SSL) enables you to ensure that any client that attempts to connect to a server during the handshake first performs server authentication. Using SSL configurations at the node, application server, and cluster scopes, you can isolate communication between servers that should not be allowed to communicate with each other over secure ports.

Before you attempt to isolate communications controlled by WebSphere Application Server, you must have a good understanding of the deployment topology and application environment. To isolate a node, application server, or cluster, you must be able to control the signers that are contained in the truststores that are associated with the SSL configuration. When the client does not contain the server signer, it cannot establish a connection to the server. By default, WebSphere uses chained certificates and each node has a unique root certificate signer. Because they the node shares the same root signer, all of the server in that node can connect to each other because they share the same root signer. However, if you use self-signed certificates, the server that created the personal certificate controls the signer, although you do have to manage the self-signed certificates. If you obtain certificates from a certificate authority (CA), you must obtain multiple CA signers because all of the servers can connect to each other if they share the same signer.

Authenticating only the server-side of a connection is not adequate protection when you need to isolate a server. Any client can obtain a signer certificate for the server and add it to its trust store. SSL client authentication must also be enabled between servers so that the server can control its connections by deciding which client certificates it can trust. For more information, see `tsec_sslclientauthinbound.dita`, which applies as well to enabling SSL client authentication at the cell level.

Isolation also requires that you use centrally managed SSL configurations for all or most endpoints in the cell. Centrally managed configurations can be scoped, unlike direct or end point configuration selection, and they enable you to create SSL configurations, key stores, and trust stores at a particular scope. Because of the inheritance hierarchy of WebSphere Application Server cells, if you select only the properties that you need for an SSL configuration, only these properties are defined at your selected scope or lower. For example, if you configure at the node scope, your configuration applies to the application server and individual end point scopes below the node scope. For more information, see `tsec_sslassocconfigscope.dita`, `tsec_sslselconfigdirect.dita`, and `tsec_sslassocconfigout.dita`

When you configure the key stores, which contain cryptographic keys, you must work at the same scope at which you define the SSL configuration and not at a higher scope. For example, if you create a key store that contains a certificate whose host name is part of the distinguished name (DN), then store that keystore in the node directory of the configuration repository. If you decide to create a certificate for the application server, then store that keystore on the application server in the application server directory.

When you configure the trust stores, which control trust decisions on the server, you must consider how much you want to isolate the application servers. You cannot isolate the application servers from the node agents or the deployment manager. However, you can configure the SOAP connector end points with the same personal certificate or to share trust. Naming persistence requires IIOP connections when they pass through the deployment manager. Because application servers always connect to the node agents when the server starts, the IIOP protocol requires that WebSphere Application Server establish trust between the application servers and the node agents.

Establishing node SSL isolation

By default, WebSphere Application Server installation uses a single chained certificate for each node so you can isolate nodes easily. A common trust store, which is located in the cell directory of the configuration repository, contains all of the signers for each node that is federated into the cell. After federation, each cell process trusts all of the other cell processes because every SSL configuration references the common trust store.

You can modify the default configuration so that each node has its own trust store, and every application server on the node trusts only the node agent that uses the same personal certificate. You must also add the signer to the node trust store so that WebSphere Application Server can establish trust with the deployment manager. To isolate the node, ensure that the following conditions are met:

- The deployment manager must initiate connections to any process
- The node agent must initiate connections to the deployment manager and its own application servers

- The application servers must initiate connections to the applications servers on the same node, to its own node agent, and the deployment manager

Figure 1 shows Node Agent A contains a key.p12 keystore and a trust.p12 trust store at the node level of the configuration repository for node A.

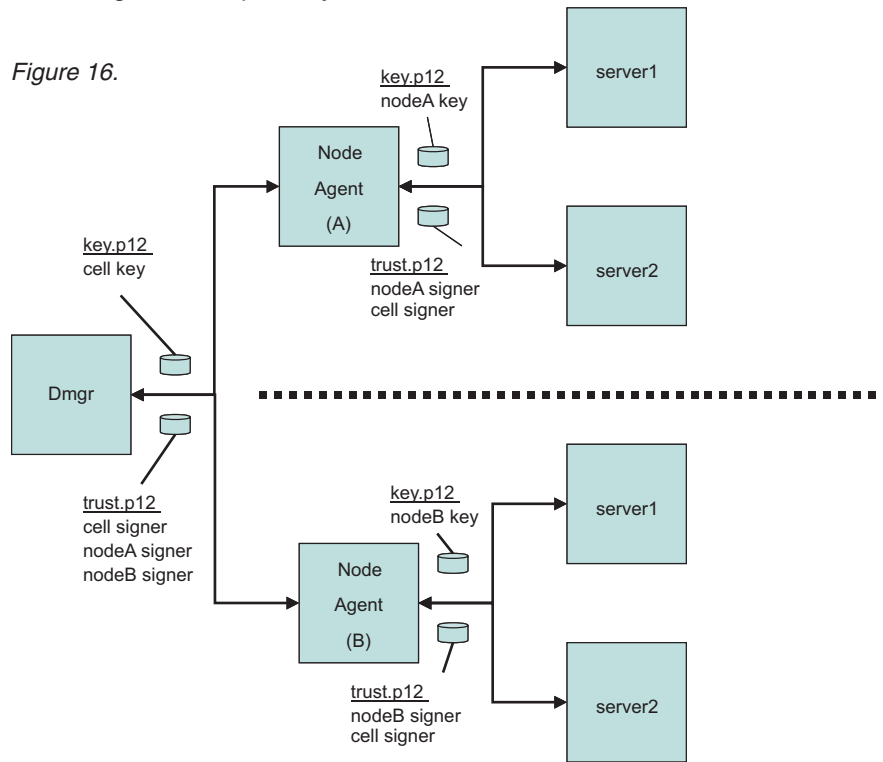


Figure 1: SSL Node Isolation

When you associate an SSL configuration with this keystore and truststore, you break the link with the cell-scoped trust store. To isolate the node completely, repeat this process for each node in the cell. WebSphere Application Server SSL configurations override the cell scope and use the node scope instead so that each process at this scope uses the SSL configuration and certificate alias that you selected at this scope. You establish proper administrative trust by ensuring that nodeA signer is in the common trust store and the cell signer is in the nodeA trust store. The same logic applies to node B as well. For more information, see `tsec_sslassocnfigscope.dita`.

Establishing application server SSL isolation

Isolating application server processes from one another is more challenging than isolating nodes. You must consider the following application design and topology conditions:

- An application server process might need to communicate with the node agent and deployment manager
- Isolating application server processes from each other might disable single sign-on capabilities for horizontal propagation

If you configure outbound SSL configurations dynamically, you can accommodate these conditions. When you define a specific outbound protocol, target host, and port for each different SSL configuration, you can override the scoped configuration.

Figure 2 shows how you might isolate an application server completely, although in practice this approach would be more complicated.

Figure 17.

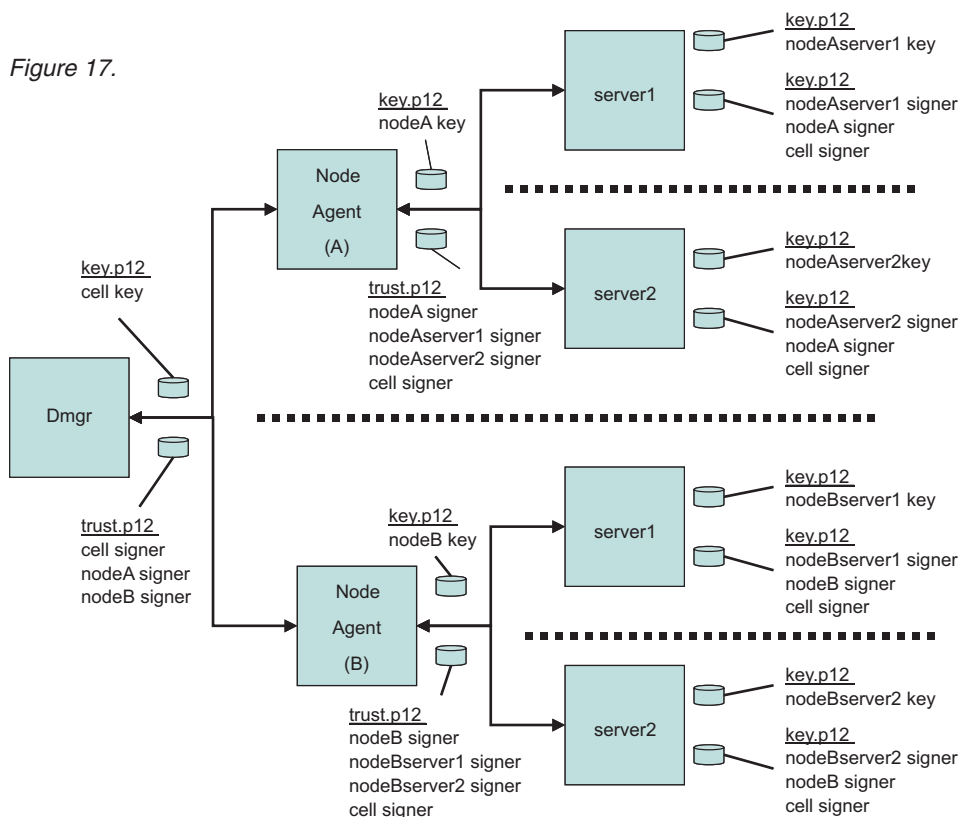


Figure 2: SSL Application Server Isolation

The dynamic configuration enables server1 on Node A to communicate with server 1 on Node B only over IIOp. The dynamic outbound rule is `IIOp,nodeBhostname,*`. For more information, see `tsec_sslssoconfigout.dita`

Establishing cluster SSL isolation

You can configure application servers into clusters instead of scoping them centrally at the node or dynamically at the server to establish cluster SSL isolation. While clustered servers can communicate with each other, application servers outside of the cluster cannot communicate with the cluster, thus isolating the clustered servers. For example, you might need to separate applications from different departments while maintaining a basic level of trust among the clustered servers. Using the dynamic outbound SSL configuration method described for servers above, you can easily extend the isolated cluster as needed.

Figure 3 shows a sample cluster configuration where cluster 1 contains a `key.p12` with its own self-signed certificate, and a `trust.p12` that is located in the `config/cells/<cellname>/clusters/<clustername>` directory.

Figure 18.




Figure 3: SSL Cluster Isolation

In the example, cluster1 might contain web applications, and cluster2 might contain EJB applications. Considering the various protocols, you decide to enable IIOp traffic between the two clusters. Your task is to define a dynamic outbound SSL configuration at the cluster1 scope with the following properties:

```
IIOP,nodeAhostname,9403|IIOP,nodeAhostname,9404|IIOP,nodeBhostname,9403|IIOP,nodeBhostname,9404
```

You must create another SSL configuration at the cluster1 scope that contains a new trust.p12 file with the cluster2 signer. Consequently, outbound IIOp requests go either to nodeAhostname ports 9403 and 9404 or to nodeBhostname ports 9403 and 9404. The IIOp SSL port numbers on these two application server processes in cluster2 identify the ports.

As you review Figure 3, notice the following features of the cluster isolation configuration:

- The trust.p12 for cluster1 contains signers that allow communications with itself (cluster1 signer), between both node agents (nodeAsigner and nodeBsigner), and with the deployment manager (cell signer).
- The trust.p12 for cluster2 contains signers that allow communications with itself (cluster2 signer), between both node agents (nodeAsigner and nodeBsigner), and with the deployment manager (cell signer).
- Node agent A and Node agent B can communicate with themselves, the deployment manager, and both clusters.

For more information, see tsec_sslssoconfigout.dita.

Although this article presents an overview of isolation methods from an SSL perspective, you must also ensure that non-SSL ports are closed or applications require the confidentiality constraint in the deployment descriptor. For example, you can set the CSiv2 inbound transport panel to require SSL and disable the channel ports that are not secure from the server ports configuration.

Also, you must enable SSL client authentication for SSL to enforce the isolation requirements on both sides of a connection. Without mutual SSL client authentication, a client can easily obtain a signer for the server programmatically and thus bypass the goal of isolation. With SSL client authentication, the server would require the client's signer for the connection to succeed. For HTTP/S protocol, the client is typically a browser, a Web Service, or a URL connection. For the IIOP/S protocol, the client is typically another application server or a Java client. WebSphere Application Server must know the clients to determine if SSL client authentication enablement is possible. Any applications that are available through a public protocol must not enable SSL client authentication because the client may fail to obtain a certificate to authenticate to the server.

Note: It is beyond the scope of this article to describe all of the factors you must consider to achieve complete isolation.

Default chained certificate configuration

When a WebSphere Application Server process starts for the first time, the Secure Sockets Layer (SSL) runtime initializes the default keystores and truststores that are specified in the SSL configuration.

Default keystore and truststore properties

WebSphere Application Server creates the `key.p12` default keystore file and the `trust.p12` default truststore file during profile creation. A default, chained certificate is also created in the **key.p12** file at this time. The root signer, or public key, of the chained certificate is extracted from the **key.p12** file and added to the **trust.p12** file. If the files do not exist during process startup, they are recreated during startup.

You can easily identify keystore and truststore defaults because of their suffixes: **DefaultKeyStore** and **DefaultTrustStore**. Also, in the SSL configuration, you must set the **fileBased** attribute to `true` so that the runtime only uses the default keystores and truststore.

On a base application server, default key and truststores are stored in the node directory of the configuration repository. For example, the default `key.p12` and `trust.p12` stores are created with the `AppSrv01` profile name, the `myhostNode01Cell` name, and the `myhostNode01` node name. The key and truststores are located in the following directories:

```
C:\WebSphere\AppServer\profiles\AppSrv01\config\cells\myhostNode01Cell
\nodes\myhostNode01\key.p12
C:\WebSphere\AppServer\profiles\AppSrv01\config\cells\myhostNode01Cell
\nodes\myhostNode01\trust.p12
```

The default password is **WebAS** for all default keystores generated by WebSphere Application Server. Change the default password after the initial configuration for a more secure environment.

Default chained certificate

The default self-signed certificate is created during profile creation for both the server and client for that profile.

You can recreate the certificates with different information simply by deleting the `*.p12` files in `/config` and `/etc`. Change the four properties below to the values you want the certificates to contain, then restart the processes. This causes the server certificate in `/config` and the client certificate in `/etc` to differ.

The certificate properties below that exist in the **ssl.client.props** file do not exist in the server configuration. However, you can use these values in the server configuration by adding them as custom security properties in the administrative console. The certificate properties appear in the **ssl.client.props** file, but do not appear in the server configuration. But, you can modify these values in the server configuration by adding them as custom properties in the administrative console.

Click **Security > Global security > Custom properties** to change the following properties:

```
com.ibm.ssl.defaultCertReqAlias=default_alias
com.ibm.ssl.defaultCertReqSubjectDN=cn=${hostname},ou=myhostNode01,ou=myhostNode01Cell,o=IBM,c=US
com.ibm.ssl.defaultCertReqDays=365
```

```

com.ibm.ssl.defaultCertReqKeySize=1024
com.ibm.ssl.rootCertSubjectDN=cn=${hostname},ou=Root Certificate, ou=myhostNode01,
ou=myhostNode01Cell,o=IBM,c=US
com.ibm.ssl.rootCertValidDays=7300
com.ibm.ssl.rootCertAlias=root
com.ibm.ssl.rootCertKeySize=1024

```

You must then delete the default chained certificate and restart the deployment manager (DMGR), node and all of the servers. If the default certificate does not exist, WebSphere Application Server automatically generates a new default certificate using the properties listed above.

If a `default_alias` value already exists, the runtime appends `_#`, where the number sign (#) is a number that increases until it is unique in the keystore. `${hostname}` is a variable that is resolved to the host name where it was originally created. The default expiration date of chained certificates is one year from their creation date.

The runtime monitors the expiration dates of chained certificates using the certificate expiration monitor. These chained certificates are automatically replaced along with any signer certificates when they are within the expiration threshold, which is typically 30 days before expiration. You can increase the default key size beyond 1024 bits only when the Java runtime environment policy files are unrestricted (that is, not exported). For more information, see “Certificate expiration monitoring” on page 622.

Default keystore and truststore configurations for new Base Application Server processes

The following sample code shows the default SSL configuration for a base application server. References to the default keystores and truststores files are highlighted.

```

<repertoire xmi:id="SSLConfig_1" alias="NodeDefaultSSLSettings"
managementScope="ManagementScope_1">
<setting xmi:id="SecureSocketLayer_1" clientAuthentication="false"
securityLevel="HIGH" enabledCiphers="" jsseProvider="IBMJSSE2" sslProtocol="SSL_TLS"
keyStore="KeyStore_1" trustStore="KeyStore_2" trustManager="TrustManager_1"
keyManager="KeyManager_1"/>
</repertoire>

```

Default keystore

In the following sample code, the keystore object that represents the default keystore is similar to the XML object.

```

<keyStores xmi:id="KeyStore_1" name="NodeDefaultKeyStore"
password="{xor}349dkckdd=" provider="IBMJCE" location="{WAS_INSTALL_ROOT}/config
/cells/myhostNode01Cell/nodes/myhostNode01/key.p12" type="PKCS12" fileBased="true"
hostList="" initializeAtStartup="true" managementScope="ManagementScope_1"/>

```

The **NodeDefaultKeyStore** keystore contains the personal certificate that represents the identity of the secure endpoint. Any keystore reference can use the `{WAS_INSTALL_ROOT}` variable, which is expanded by the runtime. The PKCS12 default keystore type is in the most interoperable format, which means that it can be imported into most browsers. The `myhostNode01Cell` password is encoded. The management scope determines which server runtime loads the keystore configuration into memory, as shown in the following code sample:

```

<managementScopes xmi:id="ManagementScope_1" scopeName="
(cells):myhostNode01Cell:(node):myhostNode01" scopeType="node"/>

```

Any configuration objects that are stored in the **security.xml** file whose management scopes are outside the current process scope are not loaded in the current process. Instead, the management scope is loaded by servers that are contained within the `myhostNode01` node. Any application server that is on the specific node can view the keystore configuration.

When you list the contents of the **key.p12** file to show the chained certificate, note that the common name (CN) of the distinguished name (DN) is the host name of the resident machine. This listing enables you to verify the host name by its URL connections. Additionally, you can verify the host name from a custom trust manager. For more information, see “Trust manager control of X.509 certificate trust decisions” on page 598.

Contents of default keystore

The following sample code shows the contents of the default **key.p12** file in a keytool list:

```
keytool -list -v -keystore c:\WebSphere\AppServer\profile\AppSrv01\profiles\config
\cells\myhostNode01Cell\nodes\myhostNode01\key.p12 -storetype PKCS12 -storepass *****
```

```
Keystore type: PKCS12
Keystore provider: IBMJCE
```

Your keystore contains 1 entry

```
Alias name: default
Creation date: Dec 31, 1969
Entry type: keyEntry
Certificate chain length: 2
Certificate[1]:
Owner: CN=myhost.austin.ibm.com, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US
Issuer: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US
Serial number: 4e48f29aaf6a6
Valid from: 2/7/08 1:03 PM until: 2/6/09 1:03 PM
Certificate fingerprints:
  MD5:  DB:FE:65:DB:40:13:F4:48:A4:CE:2F:4F:60:A5:FF:2C
  SHA1: A1:D4:DD:4B:DE:7B:45:F7:4D:AA:6A:FC:92:38:78:53:7A:99:F1:DC
Certificate[2]:
Owner: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US
Issuer: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US
Serial number: 4e48e5fd4eae3
Valid from: 2/7/08 1:03 PM until: 2/2/28 1:03 PM
Certificate fingerprints:
  MD5:  A5:9B:05:78:CF:AB:89:94:C9:2E:F1:87:34:B3:FC:75
  SHA1: 43:74:B6:C7:FA:C1:0F:19:F2:51:2B:17:60:0D:34:93:55:BF:D5:D2
```

```
*****
*****
```

The default alias name and the keyEntry entry type indicate that the private key is stored with the public key, which represents a complete personal certificate. The certificate is owned by

CN=myhost.austin.ibm.com, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US

and it is issued by the default root certificate, which is owned by

CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US

By default, the certificate is valid for one year from the date of creation.

Additionally, in some signer-exchange situations, the certificate fingerprint ensures that the sent certificate has not been modified. The fingerprint, which is a hash algorithm output for the certificate, is displayed by the WebSphere Application Server runtime during an automated signer exchange on the client side. The client fingerprint must match the fingerprint that is displayed on the server. The runtime typically uses the SHA1 hash algorithm to generate certificate fingerprints.

Default truststore

In the following sample code, the keystore object represents the default trust.p12 truststore. The truststore contains signer certificates that are necessary for making trust decisions:

```
<keyStores xmi:id="KeyStore_2" name="NodeDefaultTrustStore"
password="{xor}349dkckdd=" provider="IBMJCE" location="{WAS_INSTALL_ROOT}
/config/cells/myhostNode01Cell/nodes/myhostNode01/trust.p12" type="PKCS12"
fileBased="true" hostList="" initializeAtStartup="true" managementScope="ManagementScope_1"/>
```

Contents of default truststore

The following sample code shows the contents of the default trust.p12 truststore in a keytool listing. By default, for the sample chained certificate, the root certificate signer is included in the trust store. The root signer alias name and the trustedCertEntry entry type indicate that the certificate is the public key. The private key is not stored in this truststore. In addition, all truststores contain the default datapower certificate.

```
keytool -list -v -keystore c:\WebSphere\AppServer\profile\AppSrv01\profiles\config\cells\myhostNode01Cell1\nodes\myhostNode01\trust.p12 -storetype PKCS12 -storepass *****
```

```
Keystore type: PKCS12  
Keystore provider: IBMJCE
```

Your keystore contains 2 entries

```
Alias name: root  
Creation date: Dec 31, 1969  
Entry type: trustedCertEntry
```

```
Owner: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell1, OU=myhostNode01, O=IBM, C=US  
Issuer: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell1, OU=myhostNode01, O=IBM, C=US  
Serial number: 4e48e5fd4eae3
```

```
Valid from: 2/7/08 1:03 PM until: 2/2/28 1:03 PM
```

```
Certificate fingerprints:
```

```
MD5: A5:9B:05:78:CF:AB:89:94:C9:2E:F1:87:34:B3:FC:75
```

```
SHA1: 43:74:B6:C7:FA:C1:0F:19:F2:51:2B:17:60:0D:34:93:55:BF:D5:D2
```

```
*****  
*****
```

```
Alias name: datapower  
Creation date: Dec 31, 1969  
Entry type: trustedCertEntry
```

```
Owner: OU=Root CA, O="DataPower Technology, Inc.", C=US  
Issuer: OU=Root CA, O="DataPower Technology, Inc.", C=US  
Serial number: 0
```

```
Valid from: 6/11/03 1:23 PM until: 6/6/23 1:23 PM
```

```
Certificate fingerprints:
```

```
MD5: 18:AC:86:D1:9A:90:A2:AE:8B:28:F9:A8:75:C8:A9:DB
```

```
SHA1: A9:BA:A4:B5:BC:26:2F:5D:2A:80:93:CA:BA:F4:31:05:F2:54:14:17
```

Secure installation for client signer retrieval

Each profile in the WebSphere Application Server environment contains a unique chained certificate signed by a unique long lived root certificate that was created when the profile was created. This certificate replaces the default self-signed certificate that ships with WebSphere Application Server version 6.1 as well as the default dummy certificate that ships in releases prior to version 6.1. When a profile is federated to a deployment manager, the signer for the root signing certificate is added to the common truststore for the cell, establishing trust for all certificates signed by that root certificate.

By default, clients do not trust servers from different profiles in the WebSphere Application Server environment. That is, they do not contain the root signer for these servers. There are some things that you can do to assist in establishing this trust:

1. Enable the signer exchange prompt to except the signer during the connection attempt.
2. Run the **retrieveSigners** utility to download the signers from that system prior to making the connection.
3. Copy the trust.p12 file from the `/config/cells/<cell_name>/nodes/<node_name>` directory of the server profile to the `/etc` directory of the client. Update the SSL configuration to reflect the new file name and password, if they are different. Copying the file provides the client with a trust.p12 that contains all signers from servers in that cell. Also, you might need to perform this step for back-level clients that are still using the `DummyClientTrustFile.jks` file. In this case, you might need to change the `sas.client.props` or `soap.client.props` file to reflect the new truststore, truststore password, and truststore type (PKCS12).

For clients to perform an in-band signer exchange, you must specify the `ssl.client.props` file as a `com.ibm.SSL.ConfigURL` property in the SSL configuration. For managed clients, this is done automatically. Signers are designated either as in-band during the connection or out-of-band during runtime. You must also set the `com.ibm.ssl.enableSignerExchangePrompt` attribute to `true`.

Note: You can configure a certificate expiration monitor to replace server certificates that are about to expire. For more information about how clients can retrieve the new signer from the configuration, see “Certificate expiration monitoring” on page 622.

Using the signer exchange prompt to retrieve signers from a client

When the client does not already have a signer to connect to a process, you can enable the signer exchange prompt. The signer exchange prompt displays once for each unique certificate and for each node. After the signer for the node is added, the signer remains in the client truststore. The following sample code shows the signer exchange prompt retrieving a signer from a client:

```
C:\WASX_e0540.11\AppServer\profiles\AppSrv01\bin\serverStatus -all ADMU0116I: Tool information is being logged in file
C:\WASX_e0540.11\AppServer\profiles\AppSrv01\logs\serverStatus.log ADMU0128I: Starting tool with the
AppSrv01 profile ADMU0503I:
Retrieving server status for all servers ADMU0505I: Servers found in configuration:
ADMU0506I: Server name: dmgr
*** SSL SIGNER EXCHANGE PROMPT *** SSL signer from target host 192.168.1.5 is not found in truststore
C:\WebSphere\AppServer\profiles\AppSrv01\etc\trust.p12.
```

```
Here is the signer information (verify the digest value matches what is displayed at the server):
Subject DN:      CN=myhost.austin.ibm.com, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US
Issuer DN:      CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell, OU=myhostNode01,
O=IBM, C=US
Serial number:  2510775664686266 Expires:      Thu Feb 19 15:58:49 CST 2009
SHA-1 Digest:  2F:96:70:23:08:58:6F:66:CD:72:61:E3:46:8B:39:D4:AF:62:98:C3
MD5 Digest:    04:53:F8:20:A2:8A:6D:31:D0:1D:18:90:3D:58:B9:9D
```

```
Subject DN: CN=myhost.austin.ibm.com, OU=Root Certificate, OU=myhostNode01Cell,
OU=myhostNode01, O=IBM, C=US Issuer DN: CN=myhost.austin.ibm.com, OU=Root
Certificate, OU=myhostNode01Cell, OU=myhostNode01, O=IBM, C=US Serial number:
2510773295548841 Expires: Tue Feb 15 15:58:46 CST 2028 SHA-1 Digest:
2F:96:70:23:08:58:6F:66:CD:72:61:E3:46:8B:39:D4:AF:62:98:C3
MD5 Digest: 04:53:F8:20:A2:8A:6D:31:D0:1D:18:90:3D:58:B9:9D
Add signer to the truststore now? (y/n) y A retry of the request may need to occur. ADMU0508I:
The Deployment Manager "dmgr" is STARTED
```

To automate this process, see `rxml_retrievesigners.dita`.

When a prompt occurs to accept the signer, a socket timeout can occur and the connection might be broken. For this reason, the message **A retry of the request may need to occur.** displays after answering the prompt. The message informs the user to resubmit the request. This problem should not happen frequently, and it might be more prevalent for some protocols than others.

A retry of the request may need to occur if the socket times out while waiting for a prompt response. If the retry is required, note that the prompt will not be re-displayed if (y) is entered, which indicates the signer has already been added to the trust store.

Verify the displayed SHA-1 digest, which is the signature of the certificate that is sent by the server. If you look at the certificate on the server, verify that the same SHA-1 digest displays.

You can disable the prompt when you do not want it to display by running the **retrieveSigners** utility to retrieve all of the signers for a particular cell. You can download or upload the signers from any remote keystore to any local keystore by referencing a common truststore with this client script. For more information, see “Default chained certificate configuration” on page 614.

Using the retrieveSigners utility to download signers for a client

You can run the **retrieveSigners** utility to retrieve all of the signers from the remote keystore for a specified client keystore.

The typical remote keystore to reference is `CellDefaultTrustStore`.

The truststore contains the signers that enable the client to connect to its processes. The **retrieveSigners** utility can point to any keystore in the target configuration, within the scope of the target process, and can download the signers (certificate entries only) to any client keystore in the `ssl.client.props` file.

The following sample code shows the **retrieveSigners** utility in a deployment manager environment.

```
C:\WASX_e0540.11\AppServer\profiles\AppSrv01\bin\retrieveSigners.bat CellDefaultTrustStore
ClientDefaultTrustStore -autoAcceptBootstrapSigner CWPKI0308I: Adding signer alias
"CN=myhost.austin.ibm.com,
O=IBM, C=US" to local keystore "ClientDefaultTrustStore" with the following SHA
digest: 91:A1:A9:2D:F2:7D:70:0F:04:06:73:A3:B4:A4:9C:56:9D:A8:A3:BA CWPKI0308I:
Adding signer alias "default" to
local keystore "ClientDefaultTrustStore" with the following SHA digest:
40:20:CF:BE:B4:B2:9C:F0:96:4D:EE:E5:14:92:9E:37:8D:51:A5:47
```

Use the **-autoAcceptBootstrapSigner** option to enable WebSphere Application Server automatically to retrieve and accept the signer for administrative connections. The SHA-1 digest is printed while the signer is added so you can verify the digest after the operation is complete.

Obtaining signers for clients and servers from a previous release

Note: When a client from a release prior to version 7.0 connects to the current release, the client must obtain signers for a successful handshake. Clients using previous releases of WebSphere Application Server cannot obtain signers as easily as in the current release. You can copy the deployment manager *common* truststore to your back-level client or server, and then re-configure the SSL configuration to directly reference that truststore. This common truststore of type PKCS12 is located in the `/config/cells/<cell_name>/nodes/<node_name>` directory in the configuration repository and has a default password of WebAS.

To collect all of the signers for the cell in a single `trust.p12` keystore file, complete following steps:

1. Copy the `trust.p12` keystore file on the server and replicate it on the client. The client references the file directly from the `sas.client.props` and `soap.client.props` files that specify the SSL properties for previous releases.
2. Change the client-side keystore password so that it matches the default cell name that is associated with the copied keystore.
3. Change the default keystore type for the `trust.p12` file to PKCS12 in the client configuration.

The following two code samples show you a before and an after view of the changes to make.

Default SSL configuration of `sas.client.props` for a previous release

```
com.ibm.ssl.protocol=SSL com.ibm.ssl.keyStore=file:/// C:/SERV1_601_0208/AppServer/profiles/AppSrv01/etc/
DummyClientKeyFile.jks com.ibm.ssl.keyStorePassword={xor}CDo9Hgw\= com.ibm.ssl.keyStoreType=JKS
com.ibm.ssl.trustStore=
file:/// C:/SERV1_601_0208/AppServer/profiles/AppSrv01/etc/DummyClientTrustFile.jks
com.ibm.ssl.trustStorePassword={xor}CDo9Hgw\=
com.ibm.ssl.trustStoreType=JKS
```

SSL configuration changes that are required to common truststore file in the `/etc` directory of the client

```
com.ibm.ssl.protocol=SSL com.ibm.ssl.keyStore=file:/// C:/SERV1_601_0208/AppServer/profiles/AppSrv01/etc/
DummyClientKeyFile.jks com.ibm.ssl.keyStorePassword={xor}CDo9Hgw\= com.ibm.ssl.keyStoreType=JKS
com.ibm.ssl.trustStore=file:/// C:/SERV1_601_0208/AppServer/profiles/AppSrv01/etc/trust.p12
com.ibm.ssl.trustStorePassword=myhostNode01Cell1 com.ibm.ssl.trustStoreType=PKCS12
```

Note: You can run the **PropsFilePasswordEncoder** script, which is located in the `/bin` directory to reset the password.

You can also make these changes in the `soap.client.props` file and specify the `key.p12` file in place of the `DummyClientKeyFile.jks` file. However, you must also change the `keyStorePassword` and `keyStoreType` values to match those in the default `key.p12` file.

In releases of WebSphere Application Server prior to version 7.0, you must edit the SSL configuration on the server to replace the common truststore. The `trust.p12` file, which is used by the server, also must contain the default dummy certificate signer for connections among servers at previous release levels. You might need to manually extract the default certificate from the `DummyServerKeyFile.jks` file and then import the certificate into the `trust.p12` file that you added to the configuration.

***retrieveSigners* command:**

The `retrieveSigners` command creates a new client self-signed certificate, keystore, and SSL configuration in the `ssl.client.props` file. Using this command you can optionally extract the signer to a file.

For more information about where to run this command, read about [Using command tools](#).

Syntax

Use the following command syntax to create a new client self-signed certificate, keystore, and SSL configuration in the `ssl.client.props` file.

```
retrieveSigners <remoteKeyStoreName> <localKeyStoreName> [options]
```

The `<remoteKeyStoreName>` and `<localKeyStoreName>` parameters are required. The following optional parameters are available:

```
[-remoteAlias aliasFromRemoteStore]  
[-localAlias storeAsAlias]  
[-listRemoteKeyStoreNames] [-listLocalKeyStoreNames]  
[-autoAcceptBootstrapSigner] [-uploadSigners] [-host host]  
[-port port] [-conntype JSR160RMI|RMI|SOAP|IPC] [-user user]  
[-password password]  
[-trace] [-logfile filename]  
[-replaceLog] [-quiet] [-help]
```

Parameters

The following parameters are available for the **`retrieveSigners`** command:

`-remoteKeyStoreName`

The name of a truststore that is located in the server configuration from which to retrieve the signers. This parameter is typically the `CellDefaultTrustStore` file for a managed environment or the `NodeDefaultTrustStore` file for an unmanaged environment.

`-localKeyStoreName`

The name of the truststore that is located in the `ssl.client.props` file for the profile to which the retrieved signers is added. This parameter is typically the `ClientDefaultTrustStore` file for either a managed or unmanaged environment.

`-remoteAlias <aliasFromRemoteStore>`

Specifies one alias from the remote truststore that you want to retrieve. Otherwise, all signers from the remote truststore are retrieved.

`-localAlias <storeAsAlias>`

Determines the name of the alias stored in the local truststore. This option is only valid if you specify the `-remoteAlias` option. If you do not specify the `-localAlias` option, then the alias name from the remote truststore is used, if possible. If an alias clash occurs, then the alias name is used and has an incremented number appended to the end of it until a unique alias is found.

-listRemoteKeyStoreNames

Sends a remote request to the server to list all keystores that you can specify for the `remoteKeyStoreName` parameter. Use this command when you are unsure of the name of the remote truststore from which you want to download the signers.

-listLocalKeyStoreNames

Lists the keystores located in the `ssl.client.props` file that you can specify for the `localKeyStoreName` parameter. This truststore receives the signers from the server. Use this parameter when you are unsure of the name of the local truststore into which you want to retrieve the signers. The default name of the truststore is `ClientDefaultTrustStore` and is located in the `ssl.client.props` file.

-autoAcceptBootstrapSigner

Automatically adds a signer to make a secure connection to the server. The purpose of the option is to support automation of the command so that you do not need to accept the signer. After the signer is added to the local truststore, an SHA hash prints so that you can verify the certificate.

-uploadSigners

Converts the signer download into a signer upload. The signers from the `localKeyStoreName` parameter is sent to the `remoteKeyStoreName` parameter instead.

-host <host>

Specifies the target host from which the signers are retrieved.

-port <port>

Specifies the target administrative port to which you want to connect. You must specify the port based on the `-conntype` parameter. If the `conntype` is SOAP, the default port is 8879. This value can vary for different servers. If the `conntype` is RMI, the default port is 2809.

-conntype <JSR160RMIIPCIRMIISoap>

Determines the administrative connector type that is used for the MBean call to retrieve the signers.

Note: Eventually switch from the RMI connector to the JSR160RMI connector because support for the RMI connector is deprecated.

-user <user>

When the `-uploadSigners` flag is used, you are required to specify this option to supply the user name that is authenticated for the MBean operation. If you do not specify this parameter when the `-uploadSigners` flag is used, then you are prompted for credentials by default.

-password <password>

When the `-uploadSigners` flag is used, you are required to specify this option to supply the password that is authenticated for the MBean operation. The password goes along with the `-user` parameter.

-trace

When specified, this parameter enables tracing of the trace specification necessary to debug this component. By default, the trace is located in the `profiles/profile_name/log/retrieveSigners.log` file.

-logfile <filename>

Overrides the default trace file. By default, the trace will appear in the `profiles/profile_name/log/retrieveSigners.log` file.

-replacelog

Causes the existing trace file to be replaced when the command runs.

-quite

Suppresses most messages from printing to the console.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax for using the **retrieveSigners** command:

- The following example lists remote and local keystores:

```
retrieveSigners.bat -listRemoteKeyStoreNames -listLocalKeyStoreNames -conntype RMI -port 2809 [Windows systems]
```

```
retrieveSigners.sh -listRemoteKeyStoreNames -listLocalKeyStoreNames -conntype RMI -port 2809 [Unix systems]
```

Example output

```
CWPKI0306I: The following remote keystores exist on the specified server:
           CMSKeyStore, NodeLTPAKeys, NodeDefaultTrustStore, NodeDefaultKeyStore
CWPKI0307I: The following local keystores exist on the client:
           ClientDefaultKeyStore, ClientDefaultTrustStore
```

- The following example retrieves all signers from **NodeDefaultTrustStore**:

```
retrieveSigners.bat NodeDefault TrustStore ClientDefaultTrustStore -autoAcceptBootstrapSigner
-conntype RMI -port 2809 [Windows]
```

```
retrieveSigners.sh NodeDefault TrustStore ClientDefaultTrustStore -autoAcceptBootstrapSigner
-conntype RMI -port 2809 [Unix]
```

Example output

```
CWPKI0308I: Adding signer alias "CN=BIRKT40.austin.ibm.com, O=IBM, C=US" to
           local keystore "ClientDefaultTrustStore" with the following SHA
           digest: 40:20:CF:BE:B4:B2:9C:F0:96:4D:EE:E5:14:92:9E:37:8D:51:A5:47
```

Certificate expiration monitoring

The certificate expiration monitor administrative task is a scheduled task that cycles through all the keystores in the security configuration and reports on any certificates that are expired, certificates that fall within the expiration threshold, and certificates that fall within the pre-notification period.

Certificate expiration monitoring relies on the following definitions:

Expired certificates

Certificates are created with a finite life span. Self-signed or chained certificates that have reached the end of their life span are reported and replaced, if possible. Certificate authority (CA) signed certificates cannot be replaced but will be reported. Replacing CA-signed certificates is the responsibility of the administrator.

Certificates within the expiration threshold

There is a period of time before a certificate expires. A certificate in this period of time is one within the expiration threshold. The server replaces certificates within the expiration threshold so that the certificate does not expire and cause outages. By default the expiration threshold is 60 days, but can be configured as required.

Pre-notification period

Before a certificate falls within the expiration threshold there are warnings issued that indicate that the certificate will be replaced, when the expiration threshold date is reached. The period of time prior to the expiration threshold date is called the pre-notification period and is set at 90 days for the certificate.

The certificate expiration monitor performs the following:

1. Clears out the **NodeDefaultDeletedStore** or **DmgrDefaultDeletedStore**. This operation is performed silently without reporting that the certificates are deleted.
2. Checks the root key stores, **DmgrDefaultRootStore** or **NodeDefaultRootStore** and the **DmgrRSATokenRootStore** or **NodeRSATokenRootStore**. If any root certificates are expired, falls in the threshold period, or the pre-notification period, then the certificate is noted in the report.
3. If there are any root certificates that are expired or fall in the threshold period that root certificate is recreated using all the information used to create the original one. Any signer certificates from the original root certificate are replaced with the signers from the new root certificate.

4. If a root certificate is replaced, then all the keystores are checked to see if there are any chained certificates signed with the original root certificate. If there are, then the chain certificate is renewed (recreated with the new root certificate). Any signer certificate from the original certificate is replaced with the signer from the recreated certificate.
5. After all root keystores are processed, the rest of the keystores are checked for expired certificates, certificates in the expiration threshold, or certificates in the pre-notification period. Any certificate falling in any one of these categories is noted in the report.
6. If there are any expired certificates or certificates in the expiration threshold period and these certificates are self-signed certificates or chained certificates created by WebSphere, then they are replaced. If the chained certificates root is not in the root key store then it will be recreated as a default root certificate. Any signer certificates from the original certificate are replaced with the signer from the new certificate."
7. A report is generated and returned, written to a log file, or mailed.

The server default certificate is a chained certificate with a 365 day life span. It is signed with the default root certificate which has a 20 year life span.

You can configure this monitor task to run according to a particular schedule. The schedule produces the next start date that persists in the configuration and, when the date is reached, WebSphere Application Server starts the monitor to check all of the keystores for certificates that meet the expiration threshold. You can start the task manually to run at any time.

The following `security.xml` configuration object specifies when the monitor task starts, determines the certificate expiration threshold, and indicates whether you are notified in an e-mail using Simple Mail Transfer Protocol (SMTP) or in a message log.

```
<wsCertificateExpirationMonitor xmi:id="WSCertificateExpirationMonitor_1"
name="Certificate Expiration Monitor" daysBeforeNotification="30"
isEnabled="true" autoReplace="true" deleteOld="true"
wsNotification="WSNotification_1" wsSchedule="WSSchedule_2"
nextStartDate="1134358204849"/>
```

The expiration monitor replaces self-signed certificates and chained personal certificates that are signed by a root certificate in **DmgrDefaultRootStore** or **NodeDefaultRootStore**. Self-signed certificates are renewed using all the information that was used to create the original self-signed certificate. A chained certificate is renewed using the same root certificate that was used to sign the original certificate.

The expiration monitor automatically replaces only self-signed certificates and chained certificates that are expired or that meet the expiration threshold criteria. To replace all of the signers from the old certificate with the signer that belongs to the new certificate in all the keystores in the configuration for that cell, set the `autoReplace` attribute to `true`. When the `deleteOld` attribute is `true`, the old personal certificate and old signers also are deleted from the keystores. The `isEnabled` attribute determines whether the expiration monitor task runs based upon the `nextStartDate` attribute that is derived from the schedule. The `nextStartDate` attribute is derived from the schedule in milliseconds since 1970, and is identical to the `System.currentTimeMillis()`. If the `nextStartDate` has already passed when an expiration monitor process begins, and the expiration monitor is enabled, the task is started, but a new `nextStartDate` value is established based on the schedule.

The following sample the schedule object shows the frequency attribute as the number of days between each run of the certificate monitor.

```
<wsSchedules xmi:id="WSSchedule_2" name="ExpirationMonitorSchedule"
frequency="30" dayOfWeek="1" hour="21" minute="30"/>
```

The `dayOfWeek` attribute adjusts the schedule to run on a specified day of the week, which is always the same day regardless of whether the frequency is set to 30 or 31 days. Based on 24-hour clock, the hour and minute attributes determine when the expiration monitor is started on the specified day.

The following sample code of the notification object shows the notification configuration, which notifies you after the expiration monitor runs.

```
<wsNotifications xmi:id="WSNotification_1" name="MessageLog" logToSystemOut="true" emailList=""/
```

For expiration monitor notifications, you can select message log, e-mail using SMTP server, or both methods of notification. When you configure the e-mail option, use the format `user@domain@smtpserver`. If you do not specify an SMTP server, WebSphere Application Server defaults to the same domain as the e-mail address. For example, if you configure `joeuser@ibm.com`, WebSphere Application Server attempts to call `smtp-server.ibm.com`. To specify multiple e-mail addresses using scripting, you must add a pipe (|) character between entries. When you specify the `logToSystemOut` attribute, the expiration monitor results are sent to the message log for the environment, which is typically the `SystemOut.log` file.

The expiration monitor clears out the deleted certificates keystore. The monitor first clears out the deleted keystore. Due to the nature of the PKCS12 keystore, there must be at least one entry in the keystore so the signer certificates from the dummy key store will remain in the deleted keystore. There is no reporting on the certificate being deleted from the deleted keystore.

Note: When the expiration monitor replaces certificates, this can dynamically affect the runtime when the following configuration option is enabled:

Security > SSL certificate and key management. Under configuration settings, check the checkbox for **Dynamically update the run time when SSL configuration changes occur**.

When enabled, any certificates that are replaced causes the client SSL runtime to begin using the new certificates immediately, which in turn, flushes SSL and keystore caches and causes some ports using `SSLServerSockets` (RMI/IIOP on distributed and Admin SOAP) to restart. Restarting ports breaks existing connections. These connections can be reconnected after the port restart is completed. Endpoints using the channel framework (HTTP, BUS, RMI/IIOP on z/OS) leave existing connections unaffected but still use the new certificates for new connections.

When the dynamic change property is disabled and before the new certificates become effective, the administrator needs to recycle all processes in the entire cell after each node has the synchronized configuration. Regardless of which method is chosen, you should always check the health of your cell after the certificate expiration monitor has run (based on the schedule specified). The schedule should be set to run the certificate expiration monitor during a maintenance period so that if a restart is required after the certificate replacement, it will not cause unexpected outages.

Dynamic configuration updates

During the Secure Sockets Layer (SSL) runtime, dynamic configuration updates affect both inbound and outbound SSL endpoints. For inbound SSL endpoints, the changes that are implemented by the SSL channel are only affected by dynamic changes. For outbound SSL endpoints, all outbound connections inherit the new configuration changes.

In this release, dynamic update functionality provides you with greater flexibility and efficiency. You can change SSL configurations without restarting WebSphere Application Server for the changes to take effect.

To make dynamic changes, in the administrative console click **Security > SSL certificates and key management**, then select the **Dynamically update the runtime when SSL configuration changes occur** check box. You must save your changes and then synchronize the `security.xml` file with remote systems. A remote system must be able to confirm that `dynamicallyUpdateSSLConfig=true` is in the `security.xml` file.

The SSL runtime reloads the modified SSL configuration and creates a new `SSLEngine` for the modified connections that are associated with inbound endpoints. New outbound connections use the new configuration while existing connections continue to use the old `SSLEngine` object and are not affected.

Note: Make dynamic changes to the SSL configuration during off-peak hours. Synchronization delays can negatively affect connections when you update SSL configurations during peak hours.

You can turn on and off the `dynamicallyUpdateSSLConfig` attribute in the `security.xml` file to ensure successful updates by doing the following actions:

1. Set `dynamicallyUpdateSSLConfig=On`.
2. Save the updated configuration.
3. Synchronize the `security.xml` file with remote systems.
4. Set the `dynamicallyUpdateSSLConfig` attribute to `Off`.

You must verify that all of the nodes receive the changes before turning off the `dynamicallyUpdateSSLConfig` attribute. Test the changes in a test environment before updating the production environment.


Note: Some SSL changes, especially administrative SSL changes, can cause server outages if you fail to test them first. When a change prevents trust between two endpoints, the endpoints cannot communicate with each other. Additionally, if administrative SSL connection updates cause system outages, you might need to disable the nodes after you make corrective changes using the deployment manager. From the command line, you can manually synchronize the server to retrieve the new SSL changes, then restart the nodes.

Management scope configurations

Inbound and outbound management scopes represent opposing directions during the connection handshake process. To view inbound and outbound management scopes, use the topology tree view in the administrative console. You can define Secure Sockets Layer (SSL) configurations to distinguish the connection requirements for each direction inbound or outbound.

When expanded, the topology tree represents inbound and outbound connections for each management scope, cell, node group, node, server, cluster, and endpoint. Inbound endpoints require a server certificate. The SSL configuration specifies the server certificate for server authentication. Outbound endpoints require validated signers. Outbound endpoints connect to one or more target servers; inbound endpoints receive requests from one or more clients. The set of peer endpoints for outbound connections is typically a subset of the set of peer endpoints for inbound connections, which means you must define different requirements for inbound and outbound connections.

The following figure shows an example of two nodes: Node1 and Node2. These two nodes are isolated from one another because their SSL configurations, truststore files, and keystore files are different.



DynamicOutbound. AdminSoapSSLConfig,default ->ADMIN_SOAP
(to allow admin communications to use the same SSL configuration)

In the example of two nodes, note that Node1 cannot communicate with Node2, but each of the two nodes must be able to communicate with the deployment manager and its administrative functions. With dynamic outbound selection, you can choose an SSL configuration and a certificate alias that reference a common truststore. When a process requires the ADMIN_SOAP protocol for an outbound connection, the server uses this single SSL configuration. Because all of the scopes under the cell level inherit this configuration, all outbound connections can communicate with the deployment manager. See additional information about dynamic outbound selection of Secure Sockets Layer configurations.

Another way to accomplish this same result is to associate the SSL configuration with the ADMIN_SOAP endpoint for each individual process, deployment manager, Node1, Node2, Node1Server1, Node1Server2, Node2Server1, and Node2Server2. However, it is recommended that you use dynamic outbound selection because it is more efficient when defining a basic SSL configuration, its keystores, and its truststores at the cell scope. The example shows how to apply the node scope association, but the same principles apply for node groups, clusters, servers, and endpoints.

Note: If your topology includes clusters that span nodes or if your applications need to communicate between nodes, the configuration that is shown in the example does not work.

Certificate management using iKeyman

Starting in WebSphere Application Server Version 6.1, you can manage your certificates from the administrative console. When using versions of WebSphere Application Server prior to Version 6.1, use iKeyman for certificate management. iKeyman is a key management utility.

WebSphere Application Server certificate management requires that you define the keystores in your WebSphere Application Server configuration. With iKeyman, you need access to the keystore file only. You can read a keystore file with personal certificates and signers that is created in iKeyman. A keystore file can be read into the WebSphere Application Server configuration by using the **createKeyStore** command.

The majority of certificate management functions are the same between WebSphere Application Server and iKeyman, especially for personal certificates and signer certificates. However, certificate requests are special. The underlying behavior is different in the two certificate management schemes. Because of the different behavior, when a certificate request is generated from iKeyman, the process must be completed in iKeyman. For example, a certificate that is generated by a certificate request that originated in iKeyman must be received in iKeyman as well.

The same is true for WebSphere Application Server. For example, when a certificate is generated from a certificate request that originated in WebSphere Application Server, the certificate must be received in WebSphere Application Server.

You can perform the following certificate operations using iKeyman:

Types of certificates	Functions	Description
Personal certificates	Create a self-signed certificate	Creates a self-signed certificate and stores it in a keystore.
	List personal certificates	Lists all the personal certificates in a keystore.
	Get information about a personal certificate	Gets information about a personal certificate.
	Delete a personal certificate	Deletes a personal certificate from a keystore.
	Import a certificate	Imports a certificate from a keystore to a keystore.
	Export a certificate	Exports a certificate from a keystore to another keystore.
	Extract a certificate	Extracts the signer part of a personal certificate to a file.
Signer certificates	Receive a certificate	Reads a certificate that comes from a certificate authority (CA) into a keystore.
	Add a signer certificate	Adds a signer certificate from a file to a keystore.
	List signer certificates	Lists all the signer certificates in a keystore.
	Get information about a signer certificate	Gets information about a signer certificate.
	Delete a signer certificate	Deletes a signer certificate from a keystore.
Certificate requests	Extract a signer certificate	Extracts a signer certificate from a keystore, and stores the certificate in a file.
	Create a certificate request	Creates a certificate request that can be sent to a CA.
	List certificate requests	Lists the certificate requests in a keystore.
	Get information about a certificate request	Gets information about a certificate request.
	Delete a certificate request	Deletes a certificate request from a keystore.
	Extract a certificate request	Extracts a certificate request to a file.

Certificate management

You can manage certificate operations that involve personal certificates, signer certificates, and personal certificate requests on the administrative console.

Types of certificates

WebSphere Application Server uses the certificates that reside in keystores to establish trust for a Secure Sockets Layer (SSL) connection. Click **Security > SSL certificate and key management > Manage**

endpoint security configurations > Inbound | Outbound > *SSL_configuration_name* > Key stores and certificates, then select an existing or create a new keystore. After selecting a keystore, and depending on the type of certificate you need, choose one of the following types of certificates under Related Items:

- Personal certificate
- Signer certificate
- Certificate Authority (CA) certificates
- Personal certificate request

The following table describes the certificate operations that you can perform on the administrative console:

Types of certificates	Functions	Description
Personal certificates	Create a self-signed certificate	Creates a self-signed certificate and stores it in a keystore.
	List personal certificates	Lists all the personal certificates in a keystore.
	Get information about a personal certificate	Gets information about a personal certificate.
	Delete a personal certificate	Deletes a personal certificate from a keystore.
	Import a certificate	Imports a certificate from a keystore to a keystore.
	Export a certificate	Exports a certificate from a keystore to another keystore.
	Extract a certificate	Extracts the signer part of a personal certificate to a file.
	Exchange signer certificates	Exchange signer part of a personal certificate between key store.
	Receive a certificate	Reads a certificate that comes from a certificate authority (CA) into a keystore.
	Replace a certificate	Replaces all occurrences of a personal certificate alias in the WebSphere Application Server configuration with another certificate. Also, replaces all occurrences of the personal certificates signer with the new personal certificate signer.
Certificate authority (CA) certificates	Create a chained certificate	Creates a chained certificate and stores it in a keystore.
	Renew a certificate	Renews a certificate with a new public/private key pair and stores it in a keystore.
Signer certificates	Create CA certificate	Sends a certificate request to an external certificate authority (CA).
	Revoke CA certificate	Sends a revocation request to an external certificate authority (CA).
	Add a signer certificate	Adds a signer certificate from a file to a keystore.
	List signer certificates	Lists all the signer certificates in a keystore.
	Get information about a signer certificate	Gets information about a signer certificate.
	Delete a signer certificate	Deletes a signer certificate from a keystore.
	Extract a signer certificate	Extracts a signer certificate from a keystore, and stores the certificate in a file.
	Retrieve a signer from a port	Retrieves a signer certificate from a port, and stores it in a key store.

Types of certificates	Functions	Description
Certificate requests	Create a certificate request	Creates a certificate request that can be sent to a CA.
	List certificate requests	Lists the certificate requests in a keystore.
	Get information about a certificate request	Gets information about a certificate request.
	Delete a certificate request	Deletes a certificate request from a keystore.
	Extract a certificate request	Extracts a certificate request to a file.

Personal certificates

The following table lists the operations that you can perform on personal certificates, the AdminTask object that you can use to perform that operation, and how to navigate to the certificate on the console:

Function	AdminTask object	Administrative console
Create a self-signed certificate	createSelfSignedCertificate	Security > Secure Communications > Key store and certificates > key store > Create a Self-Signed Certificate
List personal certificates	listPersonalCertificates	Security > Secure Communications > Key store and certificates > key store > personal certificates
Get information about a personal certificate	getPersonalCertificate	Security > Secure Communications > Key store and certificates > key store > personal certificates > alias
Delete a personal certificate	deletePersonalCertificate	Security > Secure Communications > Key store and certificates > key store > personal certificates > delete
Import a certificate	importCertificate	Security > Secure Communications > Key store and certificates > key store > personal certificates > import
Export a certificate	exportCertificate	Security > Secure Communications > Key store and certificates > key store > personal certificates > export
Extract a certificate	extractCertificate	Security > Secure Communications > Key store and certificates > key store > personal certificates > extract
Exchange signer certificates	exchangeSignerCertificates	Security > Secure Communications > Key store and certificates > Exchange signers
Create a chained certificate	createChainedCertificate	Security > SSL certificate and key management > Key store and certificates > keystore name > Personal certificates. Click Create button and select Chained certificate
Renew a certificate	renewChainedCertificate	Security > SSL certificate and key management > Key store and certificates > keystore name > Personal certificates. Select a certificate. Click Renew button.

Certificate authority (CA) certificates

The following table lists the operations that you can perform on CA certificates, the AdminTask object that you can use to perform that operation, and how to navigate to the certificate on the console:

Function	AdminTask object	Administrative console
Create a CA certificate	createCACertificate	Security > Secure Communications > Key store and certificates > <i>key store</i> > Personal certificates > Create > CA-signed certificate
Revoke a CA certificate	revokeCACertificate	Security > Secure Communications > Key store and certificates > <i>key store</i> > Personal certificates <i>personal certificate</i> > Revoke

Signer certificates

The following table lists the operations that you can perform with signer certificates, the AdminTask object that you can use to perform the operation, and how to navigate to the certificate on the console:

Function	AdminTask object	Administrative console
Add a signer certificate	addSignerCertificate	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificates > Add
List signer certificates	listSignerCertificates	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificates
Get information about a signer certificate	getSignerCertificate	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificates > <i>alias</i>
Delete a signer certificate	deleteSignerCertificate	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificate > delete
Extract a signer certificate to a file	extractSignerCertificate	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificates > extract
Retrieve a signer certificate from a port	retrieveSignerFromPort	Security > Secure communications > Key store and certificates > <i>key store</i> > signer certificates > retrieve from port

Personal certificate requests

The following table lists the operations that you can perform on personal certificate requests, the AdminTask object that you can use to perform that operation, and how to navigate to the certificate request on the console:

Function	AdminTask object	Administrative console
Create a personal certificate request	createCertificateRequest	Security > Secure communications > Key store and certificates > <i>key store</i> > Personal certificate Requests > Add
List personal certificate requests	listCertificateRequests	Security > Secure communications > Key store and certificates > <i>key store</i> > Personal certificate requests
Get information about a personal certificate request	getCertificateRequest	Security > Secure communications > Key store and certificates > <i>key store</i> > Personal certificate requests > <i>alias</i>
Delete a personal certificate request	deleteCertificateRequest	Security > Secure communications > Key store and certificates > <i>key store</i> > Personal certificate requests > delete
Extract a personal certificate request to a file	extractCertificateRequest	Security > Secure communications > Key store and certificates > <i>key store</i> > Personal certificate requests > Extract

Using the retrieveSigners command to enable server to server trust

You can add a signer certificate to a server's trust.p12 file, allowing that server to securely communicate with another server. This can be done using the **retrieveSigners** command to add a signer to a server's trust.p12 file after making changes to the ssl.client.props file.

Before you begin

The server that will be communicating as a client must be identified before the server to server trust can be established. You will make change to the ssl.client.props file and run the **retrieveSigners** command on the server communicating as a client. If both servers will be acting as a client , these steps will be required for both servers.

About this task

The ssl.client.props file is setup by default to configure Secure Socket Layer (SSL) communication for clients. This makes the default behavior of the **retrieveSigners** command work on the client's trust.p12 file and key.p12 file in the *profile_root/etc* directory. You can add a signer certificate to a server's trust.p12 file, allowing that server to act as a client communicating to another server. Using the **retrieveSigners** command to add a signer to a server's trust.p12 file requires some changes to the ssl.client.props file.

1. Open the ssl.client.props file. The ssl.client.props file is located in *profile_root/properties* ditrectory.
2. Uncomment the section of ssl.client.props that starts with com.ibm.ssl.alias=AnotherSSLSettings property.
3. Uncomment the section of ssl.client.props that starts with com.ibm.ssl.trustStoreName=AnotherTrustStore property.
4. Enter the location of the trust store that the signer should be added. If you are using the server trust store for a deployment manager then it is located in *profile_root/config/cells/cell name/trust.p12*. If using the trust store for an application server, it is located in *profile_root/config/cells/cell name/nodes/node name/trust.p12*.
5. Update the remaining properties in this section with the values associated with the trust store being used. A description of the properties can be found in ssl.client.props client configuration file.
6. Optional: Uncomment and update section that starts with com.ibm.ssl.trustStoreName=AnotherKeyStore property. Most scenarios only require a signer to be added to the trust store. This example only adds a signer to the trust store, but you can also add a signer to the key store by updating the properties as you did for the trust store in steps 3 through 5.
7. Save the changes made to ssl.client.props.
8. Run the **retrieveSigners** command. For more information about where to run this command, see Using command tools

```
retrieveSigners NodeDefaultTrustStore AnotherTrustStore -host ademyers.austin.ibm.com -port 8879
```

Example output:

```
CWPKI0308I: Adding signer alias "default_1" to local keystore
           "AnotherTrustStore" with the following SHA digest:
           F4:71:97:79:3E:C1:DC:E7:9F:8F:3D:F0:A0:15:1E:D1:44:73:2C:06
```

Results

After the steps have been successfully completed. the server acting as a client has the signing certificate of the other server. This allows that server to establish a SSL connection to the other server.

Example

The example shows the modified section of the `ssl.client.props` file assuming that the server's `trust.p12` file is being used. Any trust store existing trust store can be used if the properties are provided for that trust store.

```
#-----  
com.ibm.ssl.alias=AnotherSSLSettings  
com.ibm.ssl.protocol=SSL_TLS  
com.ibm.ssl.securityLevel=HIGH  
com.ibm.ssl.trustManager=IbmX509  
com.ibm.ssl.keyManager=IbmX509  
com.ibm.ssl.contextProvider=IBMJSSE2  
com.ibm.ssl.enableSignerExchangePrompt=true  
#com.ibm.ssl.keyStoreClientAlias=default  
#com.ibm.ssl.customTrustManagers=  
#com.ibm.ssl.customKeyManager=  
#com.ibm.ssl.dynamicSelectionInfo=  
#com.ibm.ssl.enabledCipherSuites=  
  
# KeyStore information  
#com.ibm.ssl.keyStoreName=AnotherKeyStore  
#com.ibm.ssl.keyStore=${user.root}/etc/key.p12  
#com.ibm.ssl.keyStorePassword={xor}CDo9Hgw=  
#com.ibm.ssl.keyStoreType=PKCS12  
#com.ibm.ssl.keyStoreProvider=IBMJCE  
#com.ibm.ssl.keyStoreFileBased=true  
  
# TrustStore information  
com.ibm.ssl.trustStoreName=AnotherTrustStore  
com.ibm.ssl.trustStore=${user.root}/config/cells/localhostCell101/trust.p12  
com.ibm.ssl.trustStorePassword={xor}CDo9Hgw=  
com.ibm.ssl.trustStoreType=PKCS12  
com.ibm.ssl.trustStoreProvider=IBMJCE  
com.ibm.ssl.trustStoreFileBased=true
```

What to do next

After the signer has been added, edit the `ssl.client.props` file to comment out the sections that were to used to add the signer certificate.

Creating a Secure Sockets Layer configuration

Secure Sockets Layer (SSL) configurations contain the attributes that you need to control the behavior of client and server SSL endpoints. You create SSL configurations with unique names within specific management scopes on the inbound and outbound tree in the configuration topology. This task shows you how to define SSL configurations, including quality of protection and trust and key manager settings.

Before you begin

You must decide at which scope you need to define an SSL configuration, for instance, the cell, node group, node, server, cluster, or endpoint scope, from the least specific to the most specific scope. When you define an SSL configuration at the node scope, for example, only those processes within that node can load the SSL configuration; however, any processes at the endpoint in the cell can use an SSL configuration at the cell scope, which is higher in the topology.

You must also decide which scope to associate with the new SSL configuration, according to the processes that the configuration affects. For example, an SSL configuration for a hardware cryptographic device might require a keystore that is available only on a specific node, or you might need an SSL configuration for a connection to a particular SSL host and port. For more information, see “Dynamic outbound selection of Secure Sockets Layer configurations” on page 607.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations**.

2. Select an SSL configuration link on either the Inbound or Outbound tree, depending on the process you are configuring.
 - If the scope is already associated with a configuration and alias, the SSL configuration alias and certificate alias are noted in parentheses.
 - If the parenthetical information is not included, then the scope is not associated. Instead, the scope inherits the configuration properties of the first scope above it that is associated with an SSL configuration and certificate alias.

The cell scope must be associated with an SSL configuration because it is at the top of the topology and represents the default SSL configuration for the inbound or outbound connection.

3. Click **SSL configurations** under Related Items. You can view and select any of the SSL configurations that are configured at this scope. You can also view and select these configuration at every scope that is lower on the topology.
4. Click **New** to display the SSL configuration panel. You cannot select links under Additional Properties until you type a configuration name and click **Apply**.
5. Type an SSL configuration name. This field is required. The configuration name is the SSL configuration alias. Make the alias name unique within the list of SSL configuration aliases that are already created at the selected scope. The new SSL configuration uses this alias for other configuration tasks.
6. Select a truststore name from the drop-down list. A truststore name refers to a specific truststore that holds signer certificates that validate the trust of certificates sent by remote connections during an SSL handshake. If there is no truststore in the list, see “Creating a keystore configuration for a preexisting keystore file” on page 679 to create a new truststore, which is a keystore whose role is to establish trust during the connection.
7. Select a keystore name from the drop-down list. A keystore contains the personal certificates that represent a signer identity and the private key that WebSphere Application Server uses to encrypt and sign data.
 - If you change the keystore name, click **Get certificate aliases** to refresh the list of certificates from which you can choose a default alias. WebSphere Application Server uses a server alias for inbound connections and a client alias for outbound connections.
 - If there is no keystore in the list, see “Creating a keystore configuration for a preexisting keystore file” on page 679 to create a new keystore.
8. Choose a default server certificate alias for inbound connections. Select the default only when you have not specified an SSL configuration alias elsewhere and have not selected a certificate alias. A centrally managed SSL configuration tree can override the default alias. For more information, see “Central management of Secure Sockets Layer configurations” on page 608.
9. Choose a default client certificate alias for outbound connections. Select the default only when the server SSL configuration specifies an SSL client authentication.
10. Review the identified management scope for the SSL configuration. Make the management scope in this field identical to the link you selected in Step 2. If you want to change the scope, you must click a different link in the topology tree and continue at Step 3.
11. Click **Apply** if you intend to configure Additional Properties. If not, go to Step 24.
12. Click **Quality of protection (QoP) settings** under Additional Properties. QoP settings define the strength of the SSL encryption, the integrity of the signer, and the authenticity of the certificate.
13. Select a client authentication setting to establish an SSL configuration for inbound connections and for clients to send their certificates, if appropriate.
 - If you select **None**, the server does not request that a client send a certificate during the handshake.
 - If you select **Supported**, the server requests that a client send a certificate. However, if the client does not have a certificate, the handshake might still succeed.
 - If you select **Required**, the server requests that a client send a certificate. However, if the client does not have a certificate, the handshake fails.

Note: The signer certificate that represents the client must be in the truststore that you select for the SSL configuration. By default, servers within the same cell trust each other because they use the common truststore, `trust.p12`, that is located in the cell directory of the configuration repository. However, if you use keystores and truststores that you create, perform a signer exchange before you select either **Supported** or **Required**.

14. Select a protocol for the SSL handshake.
 - The default protocol, `SSL_TLS`, supports client protocols `TLSv1`, `SSLv3`, and `SSLv2`.
 - The `TLSv1` protocol supports `TLS` and `TLSv1`. The SSL server connection must support this protocol for the handshake to proceed.
 - The `SSLv3` protocol supports `SSL` and `SSLv3`. The SSL server connection must support this protocol for the handshake to proceed.

Note: Do not use the `SSLv2` protocol for the SSL server connection. Use it only when necessary on the client side.

15. Select one of the following options:
 - A predefined Java Secure Socket Extension (JSSE) provider. The `IBMJSSE2` provider is recommended for use on all platforms which support it. It is required for use by the channel framework SSL channel. When Federal Information Processing Standard (FIPS) is enabled, `IBMJSSE2` is used in combination with the `IBMJCEFIPS` crypto provider.
 - A custom JSSE provider. Type a provider name in the **Custom provider** field.
16. Select from among the following cipher suite groups:
 - **Strong:** WebSphere Application Server can perform 128-bit confidentiality algorithms for encryption and support integrity signing algorithms. However, a strong cipher suite can affect the performance of the connection.
 - **Medium:** WebSphere Application Server can perform 40-bit encryption algorithms for encryption and support integrity signing algorithms.
 - **Weak:** WebSphere Application Server can support integrity signing algorithms but not to perform encryption. Select this option with care because passwords and other sensitive information that cross the network are visible to an Internet Protocol (IP) sniffer.
 - **Custom:** you can select specific ciphers. Any time you change the ciphers that are listed from a specific cipher suite group, the group name changes to `Custom`.
17. Click **Update selected ciphers** to view a list of the available ciphers for each cipher strength.
18. Click **OK** to return to the new SSL configuration panel.
19. Click **Trust and key managers** under Additional Properties.
20. Select a default trust manager for the primary SSL handshake trust decision.
 - Choose **ibmPKIX** when you require certificate revocation list (CRL) checking using CRL distribution points in the certificates or the online certificate status protocol (OCSP).
 - Choose **ibmX509** when you do not require CRL checking but do need increased performance. You can configure a custom trust manager to perform CRL checking, if necessary.
21. Define a custom trust manager, if appropriate. You can define a custom trust manager that runs with the default trust manager you select. The custom trust manager must implement the `javax.net.ssl.X509TrustManager` interface and, optionally, the `com.ibm.wsspi.ssl.TrustManagerExtendedInfo` interface to obtain product-specific information.
 - a. Click **Security > SSL certificate and key management > Manage endpoint security configurations > SSL_configuration > Trust and key managers > Trust managers > New**.
 - b. Type a unique trust manager name.
 - c. Select the **Custom** option.
 - d. Type a class name.

- e. Click **OK**. When you return to the Trust and key managers panel, the new custom trust manager displays in the **Additional ordered trust managers** field. Use the left and right list boxes to add and remove custom trust managers.
22. Select a key manager for the SSL configuration. By default, IbmX509 is the only key manager unless you create a custom key manager.

Note: If you choose to implement your own key manager, you can affect the alias selection behavior because the key manager is responsible for selecting the certificate alias from the keystore. The custom key manager might not interpret the SSL configuration as the WebSphere Application Server key manager IbmX509 does. To define a custom key manager, click **Security > Secure communications > SSL configurations > SSL_configuration > Trust and key managers > Key managers > New**.

23. Click **OK** to save the trust and key manager settings and return to the new SSL configuration panel.
24. Click **Save** to save the new SSL configuration.

Results

Note: You can override the default trust manager when you configure at least one custom trust manager and set the `com.ibm.ssl.skipDefaultTrustManagerWhenCustomDefined` property to `true`. Click **Custom Property** on the SSL configuration panel. However, if you change the default, you leave all the trust decisions to the custom trust manager, which is not recommended for production environments. In test environments, use a dummy trust manager to avoid certificate validation. Remember that these environment are not secure.

What to do next

In this release of WebSphere Application Server, you can associate SSL configurations with protocols using one of the following methods:

- Set the SSL configuration on the thread programmatically
- Associate the SSL configuration with an outbound protocol, and target host and port. For more information, see “Associating a Secure Sockets Layer configuration dynamically with an outbound protocol and remote secure endpoint” on page 658
- Associate the SSL configuration directly using the alias. For more information, see `tsec_sslselconfigdirect.dita`
- Manage the SSL configurations centrally by associating them with SSL configuration groups or zones that are scoped for endpoints. For more information, see “Associating Secure Sockets Layer configurations centrally with inbound and outbound scopes” on page 661.

SSL certificate and key management

Use this page to configure security for Secure Socket Layer (SSL) and key management, certificates, and notifications. The SSL protocol provides secure communications between remote server processes or endpoints. SSL security can be used for establishing communications inbound to and outbound from an endpoint. To establish secure communications, a certificate and an SSL configuration must be specified for the endpoint.

To view this administrative console page, click **Security > SSL certificate and key management**.

Configuration settings

Specifies the following administrative console tasks:

- Manage endpoint security configurations
- Manage certificate expiration

Use Federal Information Processing Standard (FIPS) algorithms

Specifies the Federal Information Processing Standard (FIPS)-compliant Java cryptography engine is enabled.

- Does not affect the SSL cryptography that is performed by the application server for z/OS System Secure Sockets Layer (SSSL).
- Does not change the JSSE provider if this cell includes any Application Server versions before the application server for z/OS Version 6.0.x.

When you select the **Use the Federal Information Processing Standard (FIPS)** option, the Lightweight Third Party Authentication (LTPA) implementation uses IBMJCEFIPS. IBMJCEFIPS supports the Federal Information Processing Standard (FIPS)-approved cryptographic algorithms for Data Encryption Standard (DES), Triple DES, and Advanced Encryption Standard (AES). Although the LTPA keys are backwards compatible with prior releases of the application server, the LTPA token is not compatible with prior releases. In prior releases, the application server did not generate the LTPA token using a FIPS-approved algorithm.

The IBMJSSE2 JSSE provider does not perform cryptographic functions directly, and therefore does not need to be FIPS-approved. Instead, the IBMJSSE2 JSSE provider uses the JCE framework for cryptographic functions and uses IBMJCEFIPS when FIPS mode is enabled.

Note: The IBMJSSEFIPS provider is not supported on the HP-UX platform. However, the IBMJSSE2 provider, which uses IBMJCEFIPS, is supported on the HP-UX platform.

Default: Disabled

Dynamically update the runtime when SSL configuration changes occur

Specifies that all of the SSL-related attributes that change must be read from the configuration dynamically after they have been saved, then reused for new connections. To avoid customer impact, it is recommended that changes to production servers be made during off-peak periods.

Default: Disabled

When this option is selected, the configuration is updated each time you configure an SSL communication.

SSL configurations for selected scopes

Use this page to display Secure Socket Layer (SSL) configurations for selected scopes, such as a cell, node, server, or cluster. From this page you can navigate to configuration panels for the following: SSL configurations, dynamic inbound and outbound endpoint SSL configurations, key stores, key sets, key set groups, key managers, and trust managers.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > *ssl_configuration***.

Name

Specifies the SSL configuration scope, which is derived from the selected object in the hierarchy.

Data type: Text

Direction

Specifies the direction for which the SSLConfig applies. Inbound refers to any listener port. Outbound refers to outbound end point connections.

Data type: Text

SSL configuration

Specifies the SSL configuration that is used by requests at this scope.

Data type: Text

Update certificate alias list

Specifies the certificate aliases contained in the key store for this SSL configuration can be selected from the **Certificate alias in key store** list. You must update the certificate list after choosing a different SSL configuration alias. If you do not update the list, you will save a certificate alias that is not contained in the SSL configuration.

Manage certificates

Specifies to open the keystore panel for the key store in this SSL configuration, which enables you to manage personal certificates, signers, and certificate requests.

Certificate alias in key store

Specifies the certificate to use in the key store.

If you select **None**, the Java Secure Sockets Extension (JSSE) key manager determines which certificate is used. If multiple certificates exist in the key store, the key manager might not consistently select the same certificate.

Data type: Text

SSL configurations collection

Use this page to define a list of Secure Sockets Layer (SSL) configurations.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **SSL configurations**.

Button	Resulting action
New	The Java Secure Socket Extension (JSSE) repertoire is for Java-based SSL communications. You can define a new JSSE configuration that can be used to create an SSLContext, URLStreamHandler, SSLSocketFactory, SSLServerSocketFactory, and so on, using the com.ibm.websphere.ssl.JSSEHelper API.
Delete	Deletes an existing JSSE configuration (administrator only). Be careful that any references to the SSL configuration have been removed prior to deleting this SSL configuration.

Name

Specifies the unique name of the SSL configuration in the management scope.

SSL configuration settings

Use this page to define Secure Sockets Layer (SSL) configuration properties.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > nodes name**. Under Related items, click **SSL configurations > New**.

Name

Specifies the unique name of the SSL configuration within the management scope in which it resides. For ways to programmatically access the properties that are configured for this SSL configuration, see the `com.ibm.websphere.ssl.JSSEHelper` application programming interface (API).

Data type: Text

Trust store name

Specifies a reference to a specific trust store used by Java Secure Sockets Extension (JSSE). The trust store holds signer certificates that validate the trust of certificates sent by remote connections during an SSL handshake.

Data type: Text
Default: *selected trust store*

Key store name

Specifies a reference to a specific key store. The key store holds personal certificates that represent the identity of one side of a connection. The public key of this personal certificate is sent to the other side of the connection to establish trust during the handshake. The remote side of the connection needs the root certificate authority (CA) certificate or self-signed public key (signer) to be in the trust store to validate this personal certificate.

Data type: Text
Default: *selected key store*

Get certificate aliases

Queries the keystore for the aliases of all the personal certificates in the keystore from which to choose.

Default server certificate alias

Specifies the certificate alias used as the identity for this SSL configuration if one has not been specified elsewhere.

If you select **None**, the Java Secure Sockets Extension (JSSE) key manager determines which certificate is used. If multiple certificates exist in the key store, the key manager might not consistently select the same certificate.

Data type: Text

Default client certificate alias

Specifies the certificate alias to be used if this configuration is to be used as a client.

If you select **None**, the Java Secure Sockets Extension (JSSE) key manager determines which certificate is used. If multiple certificates exist in the key store, the key manager might not consistently select the same certificate.

Data type: Text

Management scope

Specifies the scope where this SSL configuration is visible. For example, if you choose a specific node, then the configuration is visible only on that node and on any servers that are part of that node.

Data type: Text

Certificate authority (CA) client configuration

Use this page to create, modify, and configure a certificate authority (CA) client.

To view this administrative console page, click **Security > SSL Configurations and key management** . Under Related Items, click **Certificate Authority (CA) client configurations**. Then click either the **New** button or select an existing CA client by clicking on its <client_name>.

Name

Specifies the unique name of the CA client configuration. This is the name to identify the CA client object. This name needs to be unique to the scope.

Data type: String

Implementation class

Specifies the name of the module that implements the com.ibm.wsspi.ssl.WSKPIClient interface that is used to act as a client to a CA. This implementation class connects to the CA server and performs a certificate create, revoke, or replace.

Default: String

CA server host name

Specifies the host name of the CA server, if the implementation requires a host name.

Data type: String

Port

Specifies the port where the CA server will communicate, if the implementation requires a port.

Data type: String

User name

Specifies the user Id used to connect to the CA server, if the implementation requires a user to login to the CA.

Data type: String

Password

Specifies the password for the connection to the CA server.

Data type: String

Confirm password

Confirms the password that is provided in the password field.

Data type: String

Number of times to poll

Specifies the number of times to check the CA server to see if the certificate is complete. This poll number applies to the CA that does not return certificates right away.

Default: 5

Polling interval when requesting certificates

Specifies the amount of time, in minutes, between checks to the CA server to see if the certificate is complete.

Default: 10

Custom properties

Specifies arbitrary name and value pairs of data. The name is a property key, and the value is a string value that can be used to set internal system configuration properties.

Data type: string

Certificate authority (CA) client configuration collections

Use this page to define and manage certificate authority (CA) clients or view and modify existing CA clients.

This panel allows you to create a certificate authority (CA) client object in the configuration. You can also view and modify existing CA clients. The information in the CA client object can then be used by the runtime to connect to a CA server to request, revoke, or query a certificate.

To view this administrative console page, click **Security > SSL Configurations and key management >** . Under Related Items, click **Certificate Authority (CA) client configurations**.

Button	Resulting action
New	Adds a new CA client object that can be referenced by Secure Sockets Layer (SSL) configurations.
Delete	Deletes an existing CA client object.

Name

Identifies the unique name of the CA client configuration.

Implementation class

Identifies the name of the module that implements the **com.ibm.wsspi.ssl.WSKPIClient** interface that is used to act as a client to a CA.

Management Scope

Identifies the scope where this secure sockets layer (SSL) configuration is visible.

Create chained personal certificate

A chained personal certificate is a personal certificate that is created by using another personal certificate to sign it. This chaining allows a certificate to be signed with a certificate (a root certificate) that has a long life span. Root certificates are stored in the **DmgrDefaultRootStore** or **NodeDefaultRootStore**. The server's default personal certificate is a chained certificate created when the profile is created. Chained certificates can also be created after profile creation

Before you begin

You use the administrative console to create a chained personal certificate.

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Key stores and certificates**.
3. Click a **<keystore name>** to which you want to add the chained personal certificate.
4. Under Additional Properties, click **Personal certificates** .

5. Click the **Create** button and select **Chained Certificate**. The **listCertificates** AdminTask can be used to generate the list of root certificates available to sign the certificate.

Note:

6. Fill in the following information to the General Properties section as follows:
 - Supply an alias name.
 - Select Root certificate from the pull down list.
 - Key size
 - Common name
 - Validity period
 - Organization
 - Organization Unit
 - Locality
 - State/Province
 - Zip code
 - Country or region
 -
7. Click **Apply** then **OK**.

Results

The certificate is created, signed by the root certificate specified, and stored in the keystore. Once a chained personal certificate is created, the certificate can be used by the runtime for SSL communication.

What to do next

Recovering deleted certificates

The SSL configuration contains a keystore created to hold personal certificates that were deleted from other keystores in the configuration. Perform this task to recover deleted certificates.

Before you begin

The SSL configuration contains a keystore created to hold personal certificates that were deleted from other keystores in the configuration. On a stand alone application server the keystore is called `NodeDefaultDeletedStore` and on a deployment manager the keystore is called `DmgrDefaultDeletedStore`.

When a personal certificate is deleted from a keystore using the administrative console or in a script using the `deleteCertificate` AdminTask, a copy of the certificate is stored in the `DmgrDeletedKeyStore` or `NodeDeletedKeyStore`. The personal certificate takes the alias of `<keystore>_<alias>` in the deleted keystore. If the alias name is already used in that deleted keystore a `<unique number>` is appended to the alias.

A personal certificate can be recovered from the deleted keystore by importing or exporting the personal certificate to a keystore in the configuration. To recover a personal certificate using the administrative console perform the following steps:

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Key stores and certificates**.
3. From the Keystore usages drop-down list, select "Deleted certificates keystore".
4. Click **DmgrDefaultDeletedStore** or **NodeDefaultDeletedStore**.
5. Under Additional Properties, click **Personal certificates**.

6. Select a certificate.
7. Select **Export**
8. Click **OK**.
9. Perform the following:
 - • Enter the keystore password of the deleted keystore.
 - • Enter The alias to be assigned to the certificate (in the key store that will receive the certificate).
 - • Select the 'Managed key store' radio button.
 - • Select the key store from the drop down list that will receive the certificate.
 - Click **Apply** then **OK**.

Results

Note: To recover a personal certificate you can also use the `exportCertToManagedKS AdminTask` command.

Renewing a certificate

If a personal certificate has been compromised or is about to expire, then it should be renewed. Renewing a certificate recreates the certificate with all the information from the original certificate, but with a new expiration period and public/private key pair. Only self-signed certificates and chained certificates created by WebSphere can be renewed. If the certificate used to sign the chained certificate is not in the root keystore then the default root certificate is used to renew the certificate.

Before you begin

You use the administrative console to renew the certificate.

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Key stores and certificates**.
3. Click the appropriate **<keystore name>** to which you want to add the new certificate.

Note: Only self-signed certificates and chained certificates signed with root certificates from the root keystore can be renewed.

4. Under Additional Properties, click **Personal certificates** to list the personal certificates.
5. 5. Select a personal certificate from the list.
6. 6. Click the **Renew** button.
7. Click **Apply** then **OK**.

Results

The certificate is renewed in the key store selected in the path to this panel . If the certificate is not a self-signed certificate or a chained certificate signed with a root certificate from the default root store, an error is returned.

Note: If this command is used with a CA certificate, an error occurs.

Revoke a CA certificate

If a certificate authority (CA) certificate is compromised and the servers cannot trust it anymore that CA certificate can be revoked. To revoke a CA certificate, you perform the following task.

Before you begin

You use the administrative console to replace or revoke a CA certificate.

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Key stores and certificates**.
3. Click a **<keystore name>** to which you want to add the new CA certificate.
4. Under Additional Properties, click **Personal certificates** to list the personal certificates.
5. Select a certificate to revoke (a CA certificate)
6. Click the **Revoke** button.
7. Fill in the following information to the CA certificate section.
 - Revocation password
 - Revocation reason
8. Click **Apply** then **OK**.

Results

The certificate is revoked in the key store selected in the path. If the certificate selected was not a CA certificate, then an error is returned.

What to do next

Using a CA client to create a personal certificate to be used as the default personal certificate

An external certificate authority (CA) certificate can be used as the server default personal certificate. The CA certificate can be created using a CA client.

Before you begin

What you need to have before you perform this task is as follows:

- A certificate authority (CA) to make the certificate request to.
- A module that implements the **com.ibm.wsspi.ssl.WSPKIClient** interface. This module is needed to connect to the CA server and request a certificate.

You use the administrative console to view or modify a CA client.

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Certificate Authority (CA) client configurations**. A panel displaying the existing CA clients appears.
3. Click the **New** button.
4. Enter the CA client information as required.
 - Name of the CA client.
 - The management scope (selected from the drop-down list).
 - Implementation class.
 - CA server host name.
 - User name.
 - Password.
 - Confirm of password.
 - Number of times to poll.
 - Polling interval (in minutes) when requesting certificates.
 - Custom properties.

5. Click **Apply** then **Save**.
6. Navigate to the Server default key store personal certificate. **Security > SSL configuration and certificate management > Key stores and certificates > <server_default_keystore>** . Under Additional properties, click **Personal certificates**
7. Click the **Create** button and select **CA-signed certificate**
8. Fill in the following information to the CA certificate section.
 - Revocation password
 - Confirm password.
 - Select the CA client that applies to this CA certificate.

Note: You can create a new CA client to apply to this CA authority by clicking the **New** button.

- Fill in the following information to the Request Specification section:
 - Select the radio button for Predefined request alias if you have a predefined alias.
 - If you do not have a predefined alias, fill in the following fields:
 - Type an alias name in the Alias field. The alias identifies the certificate request in the keystore.
 - Type a common name (CN) value. This value is the CN value in the certificate distinguished name (DN).
 - **Optional:** Type an organization value. This value is the O value in the certificate DN.
 - **Optional:** Select a key size value. The default key size value is 1024 bits.
 - Locality
 - **Optional:** Type the State or Province value. This value is the ST value in the certificate DN.
 - **Optional:** Type a zip code value. The zip code value is the POSTALCODE value in the certificate DN.
 - **Optional:** Type a country or region value from the list. This country value is the C= value in the certificate request DN.
 - Validity period

9. Click **Apply** then **Save**.
10. Navigate to the Server Default Key store's personal certificates **Security > SSL configuration and certificate management > Key stores and certificates > <server_default_keystore>** . Under Additional properties, click **Personal certificates**
11. Select the server default personal certificate and click the **Replace** button.
12. Select the CA certificate alias from the list of aliases.
13. Click **Apply** then **Save**.

Results

The CA certificate alias replaces the alias of the default certificate in places where it is referenced in the configuration. All signer certificates from the default certificate are replaced with the signer certificate from the CA certificate.

Create a CA certificate

Certificates can be created by a certificate authority (CA) when a CAClient object is configured to connect to the CA to create the certificate. Certificates created by a certificate authority (CA) with a CA client are tracked in the security configuration in an object called CACertificate. The certificate is stored in a keystore and a CACertificate object is added to the configuration to reference the certificate. CA certificates are personal certificates.

Before you begin

Before you begin, a CA client must be created to connect to the CA server. You then use the administrative console to create a CA certificate.

1. Click **Security > SSL certificate and key management**.
2. Under Related Items, click **Key stores and certificates**.
3. Click a **<keystore name>** to which you want to add the new CA certificate.
4. Under Additional Properties, click **Personal certificates** to create a new CA certificate in the configuration.

Note: You can also create a CA certificate by using the **requestCACertificate** AdminTask .

5. Click the **Create** button and select **CA-signed Certificate**
6. Fill in the following information to the CA certificate section.
 - Revocation password
 - Confirm password.
 - Select the CA client from the pull down list.

Note: You can create a new CA client to apply to this CA authority by clicking the **New** button.

- Fill in the following information to the Request Specification section:
 - Select the radio button for a predefined request alias if a certificate request is already created.
 - If you do not have a predefined certificate request alias, fill in the following fields:
 - a. Type an alias name in the Alias field. The alias identifies the certificate request in the keystore.
 - b. Type a common name (CN) value. This value is the CN value in the certificate distinguished name (DN).
 - c. **Optional:** Type an organization value. This value is the O value in the certificate DN.
 - d. **Optional:** Select a key size value. The default key size value is 1024 bits.
 - e. Locality
 - f. **Optional:** Type the State or Province value. This value is the ST value in the certificate DN.
 - g. **Optional:** Type a zip code value. The zip code value is the POSTALCODE value in the certificate DN.
 - h. **Optional:** Type a country or region value from the list. This country value is the C= value in the certificate request DN.
7. Click **Apply** then **OK**.

Results

The certificate is stored in the keystore selected in the path to this panel and a CACertificate configuration object is created. Once a CA certificate is created the certificate can be used by the runtime for SSL communication.

An existing certificate request can be used to create the CA certificate or a new certificate request can be created. This panel uses the requestCAClient AdminTask to create the CA certificate.

Developing the WSPKIClient interface for communicating with a certificate authority

Implementing the WSPKIClient interface enables WebSphere Application Server security to communicate with a remote certificate authority (CA).

1. Initialize the WSPKIClient method, with `init(java.util.HashMap)`.

```
public void init(java.util.HashMap initAttrs) throws WSPKIException;
```

This method is called by WebSphere Application Server runtime to set up connection information to a CA.

2. • Request a certificate with `requestCertificate(byte[], X500Principal, byte[], java.util.HashMap)`.

```
public X509Certificate[] requestCertificate(byte[] certReq,  
X500Principal SubjectDN, byte[] revocationPassword,  
java.util.HashMap customAttrs) throws WSPKIException;
```

This method is called by WebSphere Application Server runtime to connect to a CA and requests a certificate signed by the authority. A `X509Certificate[]` is returned if the requested certificate is created. If a null is returned then `queryCertificate()` is called to check if the certificate is ready. This method is used when the CA requires manual intervention to process a certificate request.

You can invoke this operation from the administrative console using the “Create a CA certificate” on page 644 task and from a client using the `requestCertificate` script, see `requestCertificate` command.

3. • Revoke a certificate with `revokeCertificate(X509Certificate[], byte[], String, java.util.HashMap)`.

```
public void revokeCertificate(X509Certificate[] cert, byte[] revocationPassword,  
String revocationReason, java.util.HashMap customAttrs) throws WSPKIException;
```

This method called by WebSphere Application Server runtime to submit a request to a CA to revoke a certificate.

You can invoke this operation from the administrative console using the `revoke CA certificate` task, “Revoke a CA certificate” on page 642, or using the `revokeCertificate` script.

4. • Query a certificate with `queryCertificate(X509Certificate[], byte[], java.util.HashMap)`.

```
public X509Certificate[] queryCertificate(byte[] certReq,  
java.util.HashMap customAttrs) throws WSPKIException;
```

This method is called by WebSphere Application Server runtime to query if certificate creation is completed on the CA. A `X509Certificate[]` is returned if certificate request is complete. A null is returned if the certificate request is pending.

You perform this operation from the administrative console using the `Query` (link to `usec_sslperscertreqs.html`) option, see “Personal certificate requests collection” on page 693 and from a client using the `queryCertificate` script, see `queryCertificate` command.

Results

the `WSPKIClient` interface for communicating with a certificate authority (CA) is implemented.

Creating a custom trust manager configuration

You can create a custom trust manager configuration at any management scope and associate the new trust manager with a Secure Sockets Layer (SSL) configuration.

Before you begin

You must develop, package, and locate a Java Archive JAR file for a custom key manager in the `was.install.root/lib/ext` directory on WebSphere Application Server. For more information, see “Example: Developing a custom trust manager for custom SSL trust decisions” on page 650.

About this task

Complete the following steps in the administrative console:

1. Decide whether you want to create the custom trust manager at the cell scope or below the cell scope at the node, server, or cluster, for example.

Note: When you create a custom trust manager at a level below the cell scope, you can associate it only with a Secure Sockets Layer (SSL) configuration at the same scope or higher. An SSL configuration at a scope lower than the trust manager does not see the trust manager configuration.

- To create a custom trust manager at the cell scope, click **Security > SSL certificate and key management > Trust managers**. Every SSL configuration in the cell can select the trust manager at the cell scope.
 - To create a custom trust manager at a scope below the cell level, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Trust managers**.
2. Click **New** to create a new custom trust manager.
 3. Type a unique trust manager name.
 4. Select the **Custom** implementation setting. The custom setting enables you to define a Java class with an implementation of the javax.net.ssl.X509TrustManager Java interface and, optionally, the com.ibm.wsspi.ssl.TrustManagerExtendedInfo WebSphere Application Server interface.

Note: The standard implementation setting applies only when the trust manager is already defined in the Java security provider list as a provider and an algorithm, which is not the case for a custom trust manager.

5. Type a class name, for example, com.ibm.test.CustomTrustManager.
6. Select one of the following actions:
 - Click **Apply**, then click **Custom properties** under Additional Properties to add custom properties to the new custom trust manager. When you are finished adding custom properties, click **OK** and **Save**, then go to the next step.
 - Click **OK** and **Save**, then go to the next step.
7. Click **SSL certificate and key management** in the page navigation at the top of the panel.
8. Select one of the following actions:
 - Click **SSL configurations** under Related Items for a cell-scoped SSL configuration.
 - Click **Manage endpoint security configurations** to select an SSL configuration at a lower scope.
9. Click the link for the existing SSL configuration that you want to associate with the new custom trust manager. You can create a new SSL configuration instead of associating the custom trust manager with an existing configuration. For more information, see “Creating a Secure Sockets Layer configuration” on page 632.
10. Click **Trust and Key managers** under Additional Properties. If the new custom trust manager is not listed in the **Additional ordered trust managers** list, verify that you selected an SSL configuration scope that is at the same level or below the scope that you selected in Step 8.
11. Click **Add**. This action adds the new trust manager to the list of custom trust managers.
12. Click **OK** and **Save**.

Results

You have created a custom trust manager configuration that references a JAR file in the install directory of WebSphere Application Server and associates it with an SSL configuration during the connection handshake.

What to do next

You can create a custom trust manager for a pure client. For more information, see “TrustManagerCommands command group for the AdminTask object” on page 1004.

Trust and key managers settings

Use this page to specify trust and key managers for the selected SSL configuration.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under **Related items** click **SSL configurations > SSL_configuration_name**. Under **Additional Properties** click **Trust and key managers**.

Note: The application server checks the default trust managers first before checking the additional ordered trust managers in descending order.

Default trust manager:

Specifies the default trust manager. The default trust manager is IbmPKIX, which can be selected when certificate revocation checks must be made using the X509Certificate CRL distribution list. The other default trust manager is IbmX509.

Data type: Text
Default: ibmPKIX

Additional ordered trust managers:

Specifies additional trust managers that are used in the order shown for this SSL configuration.

Add:

Specifies to add the selection to the **Additional ordered trust managers** right-hand list.

Remove:

Specifies to remove the selection from the **Additional ordered trust managers** right-hand list.

Key manager:

Specifies the key manager that runs for this SSL configuration.

Data type: Text
Default: IbmX509

Trust managers collection

Use this page to define the implementation settings for the trust manager. A trust manager is a class that is invoked during a Secure Sockets Layer (SSL) handshake to make trust decisions about the remote end point. A default trust manager is used to validate the signature and expiration of the certificate. Custom trust managers can be plugged in to perform an extended certificate and host name check.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under **Related items**, click **Trust managers**.

Button	Resulting action
New	Adds a new trust manager that can be selected by an SSL configuration. A trust manager is invoked during an SSL handshake and can decide whether the handshake should be accepted based on the information it knows about the remote certificate and host.
Delete	Deletes an existing trust manager. Make sure the trust manager is not referenced by any SSL configuration before you delete it.

Name:

Specifies the name of the trust manager. This name is used as a selection in the SSL configuration panel.

Class name:

Specifies a class that implements the javax.net.ssl.X509TrustManager interface. Optionally, the class can implement the com.ibm.wsspi.ssl.TrustMangerExtendedInfo interface to get extended information about the connection. The class can use the information to verify the host name and so on.

Algorithm:

Specifies the algorithm name of the trust manager that is implemented by the selected provider.

Trust managers settings

This page enables you to view and set definitions for trust manager implementation settings. A trust manager is a class that gets invoked during a Secure Sockets Layer (SSL) handshake to make trust decisions about the remote end point. A default trust manager is used to validate the signature and expiration of the certificate. Custom trust managers can be plugged in to perform an extended certificate and hostname check.

To view this administrative console page, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration** . Under Related items click **Trust managers > New** .

Name:

Specifies the name of the trust manager.

Data type:	Text
Default:	ibmX509TrustManager

Standard:

Specifies that the trust manager selection is available from a Java provider that is installed in the java.security file. This provider might be shipped by the Java Secure Sockets Extension (JSSE) or might be a custom provider that implements the javax.net.ssl.X509TrustManager interface.

Default:	Enabled
-----------------	---------

Provider:

Specifies the provider name that has an implementation of the javax.net.ssl.X509TrustManager interface. This provider is typically set to IBMJSSE2.

Enabled when **Standard** is selected.

Default	IBMJCE
----------------	--------

Algorithm:

Specifies the algorithm name of the trust manager implemented by the selected provider.

Enabled when **Standard** is selected.

Default ibmX509 or IbmPKIX
Range ibmX509, IbmPKIX

Custom:

Specifies that the trust manager selection is based on a custom implementation class that implements the javax.net.ssl.X509TrustManager interface and optionally the com.ibm.wsspi.ssl.TrustManagerExctendedInfo interface to obtain additional connection information that is not otherwise available.

Default: Disabled

Class name:

Specifies a class that implements the javax.net.ssl.X509TrustManager interface. Optionally, the class can implement the com.ibm.wsspi.ssl.TrustMangerExtendedInfo interface to get extended information about the connection. The class can use the information to verify the host name and so on.

Enabled when **Custom** is selected.

Data type: Text

Example: Developing a custom trust manager for custom SSL trust decisions

The following example is of a sample custom trust manager. The custom trust manager makes no trust decisions but instead uses the information in the X.509 certificate that it references to make decisions.

After you build and package the custom trust manager, configure it either from the ssl.client.props file for a pure client or the SSLConfiguration TrustManager link in the administrative console. See “Trust manager control of X.509 certificate trust decisions” on page 598 for more information about trust managers.

Note: This example should only be used as a sample, and is not supported.

```
import java.security.cert.X509Certificate;
import javax.net.ssl.*;
import com.ibm.wsspi.ssl.TrustManagerExctendedInfo;

public final class CustomTrustManager implements X509TrustManager,
TrustManagerExctendedInfo
{
    private static ThreadLocal threadLocStorage = new ThreadLocal();
    private java.util.Properties sslConfig = null;
    private java.util.Properties props = null;

    public CustomTrustManager()
    {
    }

    /**
     * Method called by WebSphere Application Server run time to set the target
     * host information and potentially other connection info in the future.
     * This needs to be set on ThreadLocal since the same trust manager can be
     * used by multiple connections.
     *
     * @param java.util.Map - Contains information about the connection.
     */
    public void setExtendedInfo(java.util.Map info)
    {
```

```

        threadLocStorage.set(info);
    }

    /**
     * Method called internally to retrieve information about the connection.
     *
     * @return java.util.Map - Contains information about the connection.
     */
    private java.util.Map getExtendedInfo()
    {
        return (java.util.Map) threadLocStorage.get();
    }

    /**
     * Method called by WebSphere Application Server run time to set the custom
     * properties.
     *
     * @param java.util.Properties - custom props
     */
    public void setCustomProperties(java.util.Properties customProps)
    {
        props = customProps;
    }

    /**
     * Method called internally to the custom properties set in the Trust Manager
     * configuration.
     *
     * @return java.util.Properties - information set in the configuration.
     */
    private java.util.Properties getCustomProperties()
    {
        return props;
    }

    /**
     * Method called by WebSphere Application Server runtime to set the SSL
     * configuration properties being used for this connection.
     *
     * @param java.util.Properties - contains a property for the SSL configuration.
     */
    public void setSSLConfig(java.util.Properties config)
    {
        sslConfig = config;
    }

    /**
     * Method called by TrustManager to get access to the SSL configuration for
     * this connection.
     *
     * @return java.util.Properties
     */
    public java.util.Properties getSSLConfig ()
    {
        return sslConfig;
    }

    /**
     * Method called on the server-side for establishing trust with a client.
     * See API documentation for javax.net.ssl.X509TrustManager.
     */
    public void checkClientTrusted(X509Certificate[] chain, String authType)
        throws java.security.cert.CertificateException
    {
        for (int j=0; j<chain.length; j++)
        {
            System.out.println("Client certificate information:");

```

```

        System.out.println(" Subject DN: " + chain[j].getSubjectDN());
        System.out.println(" Issuer DN: " + chain[j].getIssuerDN());
        System.out.println(" Serial number: " + chain[j].getSerialNumber());
        System.out.println("");
    }
}

/**
 * Method called on the client-side for establishing trust with a server.
 * See API documentation for javax.net.ssl.X509TrustManager.
 */
public void checkServerTrusted(X509Certificate[] chain, String authType)
    throws java.security.cert.CertificateException
{
    for (int j=0; j<chain.length; j++)
    {
        System.out.println("Server certificate information:");
        System.out.println(" Subject DN: " + chain[j].getSubjectDN());
        System.out.println(" Issuer DN: " + chain[j].getIssuerDN());
        System.out.println(" Serial number: " + chain[j].getSerialNumber());
        System.out.println("");
    }
}

/**
 * Return an array of certificate authority certificates which are trusted
 * for authenticating peers. You can return null here since the IbmX509
 * or IbmPKIX will provide a default set of issuers.
 *
 * See API documentation for javax.net.ssl.X509TrustManager.
 */
public X509Certificate[] getAcceptedIssuers()
{
    return null;
}
}

```

Related concepts

“Trust manager control of X.509 certificate trust decisions” on page 598

The role of the trust manager is to validate the Secure Sockets Layer (SSL) certificate that is sent by the peer, which includes verifying the signature and checking the expiration date of the certificate. A Java Secure Socket Extension (JSSE) trust manager determines if the remote peer can be trusted during an SSL handshake.

Related tasks

“Creating a custom trust manager configuration” on page 646

You can create a custom trust manager configuration at any management scope and associate the new trust manager with a Secure Sockets Layer (SSL) configuration.

Creating a custom key manager

You can create a custom key manager configuration at any management scope and associate the new key manager with a Secure Sockets Layer (SSL) configuration.

Before you begin

You must develop, package, and locate a Java Archive (.JAR) file for a custom key manager in the `was.install.root/lib/ext` directory on WebSphere Application Server. For more information, see `rsec_ssldevcustomkeymgr.dita`.

About this task

Complete the following steps in the administrative console:

1. Decide whether you want to create the custom key manager at the cell scope or below the cell scope at the node, server, or cluster, for example.

Note: When you create a custom key manager at a level below the cell scope, you can associate it only with a Secure Sockets Layer (SSL) configuration at the same scope or higher. An SSL configuration at a scope lower than the key manager does not see the key manager configuration.

- To create a custom key manager at the cell scope, click **Security > SSL certificate and key management > Key managers**. Every SSL configuration in the cell can select the key manager at the cell scope.
 - To create a custom key manager at a scope below the cell level, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > SSL_configuration > Key managers**.
2. Click **New** to create a new key manager.
 3. Type a unique key manager name.
 4. Select the **Custom** implementation setting. With the custom setting, you can define a Java class that has an implementation on the Java interface `javax.net.ssl.X509KeyManager` and, optionally, the `com.ibm.wsspi.ssl.KeyManagerExtendedInfo` WebSphere Application Server interface. The standard implementation setting applies only when the key manager is already defined in the Java security provider list as a provider and an algorithm, which is not the case for a custom key manager. The typical standard key manager is `algorithm = IbmX509`, `provider = IBMJSSE2`.
 5. Type a class name. For example, `com.ibm.test.CustomKeyManager`.
 6. Select one of the following actions:
 - Click **Apply**, then click **Custom properties** under Additional Properties to add custom properties to the new custom key manager. When you are finished adding custom properties, click **OK** and **Save**, then go to the next step.
 - Click **OK** and **Save**, then go to the next step.
 7. Click **SSL certificate and key management** in the page navigation at the top of the panel.
 8. Select one of the following actions:
 - Click **SSL configurations** under Related Items for a cell-scoped SSL configuration.
 - Click **Manage endpoint security configurations** to select an SSL configuration at a lower scope.
 9. Click the link for the existing SSL configuration that you want to associate with the new custom key manager. You can create a new SSL configuration instead of associating the custom key manager with an existing configuration. For more information, see “Creating a Secure Sockets Layer configuration” on page 632.
 10. Click **Trust and Key managers** under Additional Properties.
 11. Select the new custom key manager in the **Key manager** drop-down list. If the new custom key manager is not listed, verify that you selected an SSL configuration scope that is at the same level or below the scope that you selected in Step 8.
 12. Click **OK** and **Save**.

Results

You have created a custom key manager configuration that references a JAR file in the installation directory of WebSphere Application Server and associates the custom configuration with an SSL configuration during the connection handshake.

What to do next

You can create a custom key manager for a pure client. For more information, see “keyManagerCommands command group for the AdminTask object” on page 978.

Example: Developing a custom key manager for custom Secure Sockets Layer key selection

The following example is of a sample custom key manager. This simple key manager returns the configured alias if it is set using the alias properties `com.ibm.ssl.keyStoreClientAlias` or `com.ibm.ssl.keyStoreServerAlias`, depending on which side of the connection the key manager is used. The key manager defers to the JSSE default `IbmX509` key manager to select an alias if these properties are not set.

After you build and package a custom key manager, you can configure it from either the `ssl.client.props` file for a pure client or by using the `SSLConfiguration KeyManager` link in the administrative console. See “Key manager control of X.509 certificate identities” on page 600 for more information about key managers.

Because only one key manager can be configured at a time for any given Secure Sockets Layer (SSL) configuration, the certificate selections on the server side might not work as they would when the default `IbmX509` key manager is specified. When a custom key manager is configured, it is up to the owner of that key manager to ensure that the selection of the alias from the SSL configuration supplied is set properly when `chooseClientAlias` or `chooseServerAlias` are called. Look for the `com.ibm.ssl.keyStoreClientAlias` and `com.ibm.ssl.keyStoreServerAlias` SSL properties.

Note: This example should only be used as a sample, and is not supported.

```
package com.ibm.test;

import java.security.cert.X509Certificate;
import com.ibm.wsspi.ssl.KeyManagerExtendedInfo;

public final class CustomKeyManager
    implements javax.net.ssl.X509KeyManager, com.ibm.wsspi.ssl.KeyManagerExtendedInfo
{
    private java.util.Properties props = null;
    private java.security.KeyStore ks = null;
    private javax.net.ssl.X509KeyManager km = null;
    private java.util.Properties sslConfig = null;
    private String clientAlias = null;
    private String serverAlias = null;
    private int clientslotnum = 0;
    private int serverslotnum = 0;

    public CustomKeyManager()
    {
    }

    /**
     * Method called by WebSphere Application Server runtime to set the custom
     * properties.
     *
     * @param java.util.Properties - custom props
     */
    public void setCustomProperties(java.util.Properties customProps)
    {
        props = customProps;
    }

    private java.util.Properties getCustomProperties()
    {
        return props;
    }
}
```

```

/**
 * Method called by WebSphere Application Server runtime to set the SSL
 * configuration properties being used for this connection.
 *
 * @param java.util.Properties - contains a property for the SSL configuration.
 */
public void setSSLConfig(java.util.Properties config)
{
    sslConfig = config;
}

private java.util.Properties getSSLConfig()
{
    return sslConfig;
}

/**
 * Method called by WebSphere Application Server runtime to set the default
 * X509KeyManager created by the IbmX509 KeyManagerFactory using the KeyStore
 * information present in this SSL configuration. This allows some delegation
 * to the default IbmX509 KeyManager to occur.
 *
 * @param javax.net.ssl.KeyManager defaultX509KeyManager - default key manager for IbmX509
 */
public void setDefaultX509KeyManager(javax.net.ssl.X509KeyManager defaultX509KeyManager)
{
    km = defaultX509KeyManager;
}

public javax.net.ssl.X509KeyManager getDefaultX509KeyManager()
{
    return km;
}

/**
 * Method called by WebSphere Application Server runtime to set the SSL
 * KeyStore used for this connection.
 *
 * @param java.security.KeyStore - the KeyStore currently configured
 */
public void setKeyStore(java.security.KeyStore keyStore)
{
    ks = keyStore;
}

public java.security.KeyStore getKeyStore()
{
    return ks;
}

/**
 * Method called by custom code to set the server alias.
 *
 * @param String - the server alias to use
 */
public void setKeyStoreServerAlias(String alias)
{
    serverAlias = alias;
}

```



```

private String getKeyStoreServerAlias()
{
    return serverAlias;
}

/**
 * Method called by custom code to set the client alias.
 *
 * @param String - the client alias to use
 */
public void setKeyStoreClientAlias(String alias)
{
    clientAlias = alias;
}

private String getKeyStoreClientAlias()
{
    return clientAlias;
}

/**
 * Method called by custom code to set the client alias and slot (if necessary).
 *
 * @param String - the client alias to use
 * @param int - the slot to use (for hardware)
 */
public void setClientAlias(String alias, int slotnum) throws Exception
{
    if ( !ks.containsAlias(alias))
    {
        throw new IllegalArgumentException ( "Client alias " + alias + "
            not found in keystore." );
    }
    this.clientAlias = alias;
    this.clientslotnum = slotnum;
}

/**
 * Method called by custom code to set the server alias and slot (if necessary).
 *
 * @param String - the server alias to use
 * @param int - the slot to use (for hardware)
 */
public void setServerAlias(String alias, int slotnum) throws Exception
{
    if ( ! ks.containsAlias(alias))
    {
        throw new IllegalArgumentException ( "Server alias " + alias + "
            not found in keystore." );
    }
    this.serverAlias = alias;
    this.serverslotnum = slotnum;
}

/**
 * Method called by JSSE runtime to when an alias is needed for a client
 * connection where a client certificate is required.
 *
 * @param String keyType

```

```

* @param Principal[] issuers
* @param java.net.Socket socket (not always present)
*/
public String chooseClientAlias(String[] keyType, java.security.Principal[]
issuers, java.net.Socket socket)
{
    if (clientAlias != null && !clientAlias.equals(""))
    {
        String[] list = km.getClientAliases(keyType[0], issuers);
        String aliases = "";

        if (list != null)
        {
            boolean found=false;
            for (int i=0; i<list.length; i++)
            {
                aliases += list[i] + " ";
                if (clientAlias.equalsIgnoreCase(list[i]))
                    found=true;
            }

            if (found)
            {
                return clientAlias;
            }
        }
    }

    // client alias not found, let the default key manager choose.
    String[] keyArray = new String [] {keyType[0]};
    String alias = km.chooseClientAlias(keyArray, issuers, null);
    return alias.toLowerCase();
}

/**
* Method called by JSSE runtime to when an alias is needed for a server
* connection to provide the server identity.
*
* @param String[] keyType
* @param Principal[] issuers
* @param java.net.Socket socket (not always present)
*/
public String chooseServerAlias(String keyType, java.security.Principal[]
issuers, java.net.Socket socket)
{
    if (serverAlias != null && !serverAlias.equals(""))
    {
        // get the list of aliases in the keystore from the default key manager
        String[] list = km.getServerAliases(keyType, issuers);
        String aliases = "";

        if (list != null)
        {
            boolean found=false;
            for (int i=0; i<list.length; i++)
            {
                aliases += list[i] + " ";
                if (serverAlias.equalsIgnoreCase(list[i]))
                    found = true;
            }
        }
    }
}

```

```

        }
        if (found)
        {
            return serverAlias;
        }
    }
}

// specified alias not found, let the default key manager choose.
String alias = km.chooseServerAlias(keyType, issuers, null);
return alias.toLowerCase();
}

public String[] getClientAliases(String keyType, java.security.Principal[] issuers)
{
    return km.getClientAliases(keyType, issuers);
}

public String[] getServerAliases(String keyType, java.security.Principal[] issuers)
{
    return km.getServerAliases(keyType, issuers);
}

public java.security.PrivateKey getPrivateKey(String s)
{
    return km.getPrivateKey(s);
}

public java.security.cert.X509Certificate[] getCertificateChain(String s)
{
    return km.getCertificateChain(s);
}

public javax.net.ssl.X509KeyManager getX509KeyManager()
{
    return km;
}
}

```

Associating a Secure Sockets Layer configuration dynamically with an outbound protocol and remote secure endpoint

After you create a Secure Sockets Layer (SSL) configuration, you must associate a secure outbound management scope with the new configuration. In this release, you can associate one SSL configuration with one remote secure endpoint and a different SSL configuration to another remote secure endpoint. Both endpoints can use the same outbound protocol, if appropriate. This task describes how to create the association dynamically.

Before you begin

Dynamic outbound selection requires that you provide only the outbound protocol name, the target host, and the target port so that WebSphere Application Server can make a connection between the SSL configuration and the outbound protocol or remote secure endpoint. The dynamic outbound selection method takes precedence over other selection methods, such as central management and direct selection, but is second to the programmatic method, that is, setting an SSL configuration on the running thread. For more information about the selection types and precedence rules, see “Secure communications using Secure Sockets Layer” on page 589.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > Outbound**.
2. Select the management scope that you want to associate with an SSL configuration on the topology tree.
3. Under Related Items, click **Dynamic outbound endpoint SSL configurations**. The default dynamic outbound configuration name, the target protocol, host, and port connection information, and the SSL configuration name display.
4. Click **New** to create a new dynamic outbound configuration.
5. Type a dynamic outbound configuration name. Use a name that is descriptive of the purpose of the dynamic selection configuration.
6. Optionally, type a dynamic selection configuration description.
7. Type the connection information that you want to associate with the configuration that is displayed in the SSL configuration drop-down list. The connection information must be in the format *protocol name, target host, target port*. You can substitute an asterisk (*) for any value, as in the following examples:
 - *,*,443
 - *,www.ibm.com,443
 - HTTP,.austin.ibm.com,*where 443 is a port, www.ibm.com is a host, HTTP is a protocol, and .austin.ibm.com is a target host. You can add multiple connections, but each additional connection can affect outbound performance.
8. Click **Add** to add the new connection to the set of SSL configuration connections. To remove a connection, select it and click **Remove**.
9. Select an SSL configuration from the list.
10. Click **Get certificate aliases** to refresh the certificate aliases that are contained in the associated key store.
11. Choose a certificate alias from the list.
12. Click **OK** and **Save**.

Results

WebSphere Application Server is ready to connect one or more SSL configurations to one or more remote secure endpoints.

What to do next

You can return to the outbound tree and select another management scope to associate with the same or a new outbound configuration.

Programmatically specifying an outbound SSL configuration using JSSEHelper API

WebSphere Application Server provides a way to specify programmatically which Secure Sockets Layer (SSL) configurations to use prior to making an outbound connection. The `com.ibm.websphere.ssl.JSSEHelper` interface provides a complete set of application programming interfaces (APIs) for handling SSL configurations.

About this task

Perform the following steps for your application when using the JSSEHelper API to establish an SSL properties object on the thread for use by the runtime. Some of these APIs have Java 2 Security permission requirements. See the JSSEHelper API documentation for more information about the permissions required by your application.

Select the approach that best fits your connection situation when you specify programmatically which Secure Sockets Layer (SSL) configurations to use prior to making an outbound connection.

1. Obtain an instance of the JSSEHelper API.

```
com.ibm.websphere.ssl.JSSEHelper jsseHelper = com.ibm.websphere.ssl.JSSEHelper.getInstance();
```

2. Obtain SSL properties from the WebSphere Application Server configuration or use those provided by your application. Use one of the following options.

- By direction selection of an alias name, within the same management scope or higher as in the following example:

```
try
{ String alias = "NodeAServer1SSLSettings";
  // As specified in the WebSphere SSL configuration Properties
  sslProps = jsseHelper.getProperties(alias); }
catch (com.ibm.websphere.ssl.SSLException e)
{ e.printStackTrace(); // handle exception }
```

- By using the `getProperties` API for programmatic, direction, dynamic outbound, or management scope selection (based on precedence rules and inheritance). The SSL runtime uses the `getProperties` API to determine which SSL configuration to use for a particular protocol. This decision is based on both the input (`sslAlias` and `connectionInfo`) and the management scope from which the property is called. The `getProperties` API makes decisions in the following order:
 - a. The API checks the thread to see if properties already exist.
 - b. The API checks for a dynamic outbound configuration that matches the `ENDPOINT_NAME`, `REMOTE_HOST`, and or `REMOTE_PORT`.
 - c. The API checks to see if the optional `sslAlias` property is specified. You can configure any protocol as direct or centrally managed. When a protocol is configured as direct, the `sslAlias` parameter is `null`. When a protocol is configured as centrally managed, the `sslAlias` parameter is also `null`.
 - d. If no selection has been made, the API chooses the dynamic outbound configuration based on the management scope it was called from. If the dynamic outbound configuration is not defined in the same scope, it then searches the hierarchy to locate one.

The last choice is the cell-scoped SSL configuration (in Network Deployment) or the node-scoped SSL configuration (in Base Application Server). The `com.ibm.websphere.ssl.SSLConfigChangeListener` parameter is notified when the SSL configuration that is chosen by a call to the `getProperties` API changes. The protocol can then call the API again to obtain the new properties as in the following example:

```
try { String sslAlias = null;
  // The sslAlias is not specified directly at this time. String host = "myhost.austin.ibm.com";
  // the target host String port = "443";
  // the target port HashMap connectionInfo = new HashMap();
  connectionInfo.put(JSSEHelper.CONNECTION_INFO_DIRECTION, JSSEHelper.DIRECTION_OUTBOUND);
  connectionInfo.put(JSSEHelper.CONNECTION_INFO_REMOTE_HOST, host);
  connectionInfo.put(JSSEHelper.CONNECTION_INFO_REMOTE_PORT, Integer.toString(port));
  connectionInfo.put(JSSEHelper.CONNECTION_INFO_ENDPOINT_NAME, JSSEHelper.ENDPOINT_IIOPI);
  java.util.Properties props = jsseHelper.getProperties(sslAlias, connectionInfo, null); }
catch (com.ibm.websphere.ssl.SSLException e)
{ e.printStackTrace(); // handle exception }
```

- By creating your own SSL properties and then passing them to the runtime, as in the following example:

```

try {
    // This is the recommended "minimum" set of SSL properties. The trustStore can
    // be the same as the keyStore. Properties sslProps = new Properties();
    sslProps.setProperty("com.ibm.ssl.trustStore", "some value");
    sslProps.setProperty("com.ibm.ssl.trustStorePassword", "some value");
    sslProps.setProperty("com.ibm.ssl.trustStoreType", "some value");
    sslProps.setProperty("com.ibm.ssl.keyStore", "some value");
    sslProps.setProperty("com.ibm.ssl.keyStorePassword", "some value");
    sslProps.setProperty("com.ibm.ssl.keyStoreType", "some value");
    jsseHelper.setSSLPropertiesOnThread(sslProps); }
catch (com.ibm.websphere.ssl.SSLException e)
{ e.printStackTrace(); // handle exception }

```

3. Use the `JSSEHelper.setSSLPropertiesOnThread(props)` API to set the `Properties` object on the thread so that the runtime picks it up and uses the same `JSSEHelper.getProperties` API. You can also obtain properties from the thread after they are set with the `jsseHelper.getSSLPropertiesOnThread()` API, as in the following example:

```

try
{ Properties sslProps = jsseHelper.getProperties(null, connectionInfo, null);
  jsseHelper.setSSLPropertiesOnThread(sslProps); }
catch (com.ibm.websphere.ssl.SSLException e)
{ e.printStackTrace(); // handle exception }

```

4. When the connection is completed, you must clear the SSL properties from the thread by passing the `null` value to the `setPropertiesOnThread` API.

```

try
{ jsseHelper.setSSLPropertiesOnThread(null); }
catch (com.ibm.websphere.ssl.SSLException e)
{ e.printStackTrace(); // handle exception }

```

Associating Secure Sockets Layer configurations centrally with inbound and outbound scopes

After you create a Secure Sockets Layer (SSL) configuration, you must associate a secure inbound or outbound management scope with the new configuration. You can manage the association centrally so that you can easily make changes that affect all the scopes that are lower on the topology and that are associated with the configuration. Beginning with WebSphere Application Server version 6.1, the recommended and the default configuration method is centrally managed SSL configurations.

Before you begin

You can simplify the number of associations that you need to make for an SSL configuration by associating the configuration with the highest level management scope requiring a unique configuration. SSL configuration associations manifest inheritance behaviors. Because of the inheritance behaviors, all of the scopes that are lower on the topology inherit this SSL configuration. For example, an association you make at the cell level affects nodes, servers, clusters, and endpoints. For more information, see “Central management of Secure Sockets Layer configurations” on page 608.

A precedence rule determines which SSL configuration association is used at a particular scope. The highest precedence is given to endpoints on the topology. If you establish an association at the endpoint, this association overrides any prior association that you made higher up on the management scope topology.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management**.
2. Select the **Dynamically update the runtime when SSL configuration changes** check box if you want changes that you make to an existing SSL configuration to occur dynamically. All outbound SSL communications honor the dynamic SSL changes. Protocols that do not use the channel frameworks

SSL channel for inbound communications, including Object Request Broker (ORB) and administrative SOAP protocols, do not honor dynamic updates. For more information, see “Dynamic configuration updates” on page 624.

3. Click **Manage endpoint security configurations**.
4. Select either the inbound or the outbound tree. After finishing the selected tree, you can return to this step to repeat the following steps for the other tree.
5. Click the link for the selected cell, node, node group, server, cluster, or endpoint on the topology tree. If the scope already has an associated SSL configuration and alias, these objects display in parentheses immediately following the scope name, for example: Node01(NodeDefaultSSLSettings,default). If the deployment manager has federated a node, the node scope SSL configuration overrides the cell scope configuration above it in the topology.
6. Decide whether to override the inherited values that display in the read-only fields. Read-only fields include the management scope name, the direction, and the inherited SSL configuration name and certificate alias.
 - If you are satisfied with these values, do not override them.
 - If you want to override the inherited values, select the **Override inherited values** check box.
7. Select an SSL configuration from the list.
8. Click **Update certificate alias list**. The certificate alias list comes from the key store that is referenced by the new SSL configuration.
9. Click **Manage certificates** if you want to manage the personal certificates that are contained in the key store that is referenced in the SSL configuration.
10. Click **Update certificate alias list** to refresh the list of aliases.
11. Select a certificate alias in the key store to represent the identity of the endpoint.
12. Click **OK** to save your changes.
13. Click **Manage endpoint security configurations and trust zones** to return to the topology tree.
14. Configure the opposite direction on the topology tree using the steps in this task. You can also select additional scopes to associate with the SSL configuration, as needed.

Results

Each SSL configuration at the selected scope and at scopes beneath it on the topology tree have the same SSL configuration properties. The following SSL configuration methods override the centrally managed configurations that you associate in the tree view:

- Direct selection at the endpoint
- Dynamic outbound SSL configuration associations
- Programmatic specifications

What to do next

At any management scope, you can configure the following objects: dynamic outbound endpoint SSL configurations, key stores, key sets, key set groups, key managers, and trust managers. Like SSL configurations, these objects are scoped automatically so that they are not visible higher up in the tree nor are they loaded during runtime by processes that are higher up in the tree.

Selecting an SSL configuration alias directly from an endpoint configuration

You can associate a secure outbound endpoint with a new Secure Sockets Layer (SSL) configuration directly. If you are migrating from a release prior to version 6.1, WebSphere Application Server still supports configurations that were selected directly at an endpoint. Direct selection always overrides centrally managed configurations and preserves migrated configurations.

About this task

Select an SSL configuration alias directly at the following endpoints:

- **Security > Global security > RMI/IIOP security > CSiv2 outbound transport**
- **Security > Global security > RMI/IIOP security > CSiv2 inbound transport**
- **System administration > Deployment manager > Transport Chain > WCInboundAdminSecure > SSL inbound channel (SSL_1)**
- **System administration > Deployment manager > Administration Services > JMX connectors > SOAPConnector > Custom Properties > sslConfig**
- **System administration > Node agents > nodeagent > Administration Services > JMX connectors > SOAPConnector > Custom Properties > sslConfig**
- **Servers > Application servers > server1 > Messaging engine inbound transports > InboundSecureMessaging > SSL inbound channel (SIB_SSL_JFAP)**
- **Servers > Application servers > server1 > WebSphere MQ link inbound transports > InboundSecureMQLink > SSL inbound channel (SIB_SSL_MQFAP)**
- **Servers > Application servers > server1 > SIP Container Settings > SIP container transport chains > SIPInboundDefaultSecure > SSL inbound channel (SSL_5)**
- **Servers > Application servers > server1 > Web Container Settings > Web container transport chains > WCInboundAdminSecure > SSL inbound channel (SSL_1)**
- **Servers > Application servers > server1 > Web Container Settings > Web container transport chains > WCInboundDefaultSecure > SSL inbound channel (SSL_2)**

Note: The central management of SSL configurations can be a more efficient strategy because multiple configurations can be contained within a single SSLConfigGroup. If you need to convert configuration references that are already directly managed to centrally managed configurations, modify each endpoint individually. Use the **AdminConfig.modify** command to set the `sslConfigAlias` value to an empty string (""). Below is an example of doing this:

- Using Jacl:

```
set s1 [$AdminConfig getid /Cell:mycell/Node:mynode/Server:server1/]
set sslChannel [lindex [$AdminConfig list SSLInboundChannel $s1] 0]
$AdminConfig modify $sslChannel [list[list sslConfigAlias ""]]
```

For more information on using this command, see [Configuring processes using scripting](#)

For more information on specific wsadmin commands that affect a repertoire as opposed to individual endpoints, see the [SSLConfigGroupCommands](#) group for the AdminTask here, "SSLConfigGroupCommands group for the AdminTask object" on page 1000.

Complete the following steps in the administrative console:

Note: These steps provide an example to follow when you directly select any of the endpoints listed above.

1. Click **Security > Global security > RMI/IIOP security > CSiv2 outbound transport**.
2. Click **Use specific SSL alias**. When you identify a specific SSL alias, you override the centrally managed scope associations.
3. Select an SSL configuration alias from the drop-down list.
4. Click **OK**.
5. Repeat these steps for additional protocols or endpoints, if desired.

Results

By associating the endpoint directly, you have overridden a centrally managed SSL configuration.

What to do next

If you decide to use management scopes instead of endpoints to associate an SSL configuration, follow the steps above, but click **Centrally managed** instead of **Use specific SSL alias**, then click **Manage endpoint security configurations**. The console is redirected to **Security > SSL certificate and key management > Manage endpoint security configurations**.

Enabling Secure Sockets Layer client authentication for a specific inbound endpoint

When you establish a Secure Sockets Layer (SSL) configuration, you can enable client authentication for a specific inbound endpoint.

Before you begin

The endpoint configuration must already exist in the SSL topology.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > Inbound > SSL_configuration**. If you want to enable SSL client authentication for all processes, define an SSL configuration for the new endpoint at the node or cell level so that it is visible to all processes on the same node or on the entire cell. For more information, see “Creating a Secure Sockets Layer configuration” on page 632.
2. Select **Override inherited values**. The SSL configuration is used for the current scope and any lower scopes that have not already designated an SSL configuration. This field displays for server and node groups within the object hierarchy and does not display for the top-level node or cell.
3. Select an SSL configuration from the drop-down list.
4. Click **Update certificate alias list**.
5. Select a **Certificate alias** from the drop-down list.
6. Click **OK** to save the configuration.

Results

You can repeat the previous steps for each endpoint that uses the same SSL configuration to enable client authentication for the inbound endpoints.

What to do next

CSlv2 Protocol Exception:

The Common Secure Interoperability Version 2 (CSlv2) secure endpoints, used for Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) security, cannot override inherited values. While the rest of the SSL properties are effective for CSlv2 when they are selected at the centrally-managed Secure Communications panel, the client authentication selection is controlled by the CSlv2 protocol configuration.

To enable SSL client certificate authentication for the CSlv2 protocol, you must use the CSlv2 inbound and outbound authentication panels. For SSL client authentication to occur between two servers, you must enable (support or require) SSL client certificate authentication for both the inbound and the outbound policies.

WebSphere Application Server can either request (support) clients to provide signer certificates for the SSL handshake, or the server can require clients to provide a valid signer certificate for the SSL

handshake, which is a more secure method. However, when the server requires certificates, the server must obtain a signer for each client that connects to the server, which involves more server-side management.

The client certificate should not be used for the identity when it is used from server-to-server. However, when a pure client sends the client certificate it is used for the identity unless a message level identity is specified, such as a user ID or a password.

Do the following to enable client certificate authentication for the CSiv2 protocol for server-to-server:

1. **Click Security > Global security.**
2. Expand the **RMI/IIOP** security section.
3. Click **CSiv2 inbound authentication.**
4. Under Client authentication, select either **supported** or **required**. When you select required, only one SSL port is opened (CSV2_SSL_MUTUALAUTH_LISTENER_ADDRESS). When you select supported, two SSL ports are opened (both CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS and CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS).
If there are two ports, the client can select either based on the security configuration policy of the port.
5. Click **OK** to save.
6. If you want server-to-server SSL client authentication, then complete the remaining steps. If you don't complete the remaining steps, only pure clients are enabled to send client certificates.
7. Expand the **RMI/IIOP** security section.
8. Click **CSiv2 outbound authentication.**
9. Under Client authentication, select either **supported** or **required**.

The SSL configuration for the inbound secure endpoints for which you enable SSL client certificate authentication must have the signer certificate from any client that attempts to open a connection to that inbound secure endpoint. You must collect those signers and then add them to the trust store associated with the inbound secure endpoints SSL configuration.

Manage endpoint security configurations

Use this page to select a Secure Socket Layer (SSL) configuration from the Local Topology hierarchy, which includes cells, nodes, node groups, servers, and clusters.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations**.

Local topology:

The Local topology represents the hierarchy of nodes, node groups, clusters, servers, and end points within the cell that comprise a centralized SSL configuration.

The topology acts as a hierarchical tree in terms of inheritance. For example, if an SSL configuration has been associated with a specific node, then all servers within that node will inherit that SSL configuration selection, provided the servers are not associated with an SSL configuration at the server scope. Centralized management of SSL is the default configuration; however, it can be overridden at various locations to directly select a specific SSL alias as in previous releases for backwards compatibility.

Scope	Description
Inbound/Outbound	Specifies the topology tree in terms of connection direction. For example, the inbound tree represents all server endpoints that receive connections at the various servers within the cell. The outbound tree represents the client side of connections from the various servers within the cell.

Scope	Description
Nodes	Specifies the nodes that are part of the cell. The list of nodes is updated anytime a node gets federated into the cell.
Servers	Specifies the servers that are part of a specific node. You can enable a specific server to have an SSL configuration associated with it so that resources within the same server can use the associated SSL configuration.
Clusters	Specifies the clusters that are part of the cell. When an SSL configuration is associated with a cluster, all servers within the cluster will use the same SSL configuration unless specified at a lower level in the topology.
Nodegroups	Specifies the node groups that are part of the cell. When an SSL configuration is associated with a node group, all nodes within that node group may use the same SSL configuration unless one is specified at a lower scope in the topology or the specific end point has chosen a direct alias reference.
Secure port and transport	Specifies an endpoint name to associate with an SSL configuration when more specific SSL settings are needed at this level. You could select an alias directly at the endpoint panel; however, when you use Secure port and transport , you can maintain more centralized control of the SSL configuration and make changes more easily.

Dynamic inbound and outbound endpoint SSL configurations collection

Use this page to manage dynamic endpoint Secure Sockets Layer (SSL) configurations, which represent associations between Secure Socket Layer (SSL) configurations and their target protocol, host, and port.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > *ssl_configuration***. Under Related items, click **Dynamic outbound endpoint SSL configurations**.

When an outbound connection is attempted, this association is checked ahead of the SSL configuration scope association. Based on the target protocol, host, port, the outbound SSL configuration used can be different from the default specified in the SSL scope configuration.

Button	Resulting action
New	Adds a new dynamic outbound selection criteria. The outbound connection selects an SSL configuration based upon connection information, including DNS host name and domain, port, and protocol type. When an outbound connection is being made, the dynamic outbound selection criteria are queried for a match, and if found the SSL configuration associated is used.
Delete	Deletes an existing dynamic outbound endpoint SSL configuration.

Name:

Specifies the unique name of the dynamic endpoint configuration.

Connection information:

Specifies the set of target protocol, host, port for the outbound request in the form *protocol,host,port*.

SSL Configuration:

Specifies the SSL configuration that is used by requests at this scope when a match occurs for the given selection criteria.

Dynamic outbound endpoint SSL configuration settings

Use this page to set properties for dynamic outbound endpoint SSL configurations, which represent associations between SSL configurations and their target protocol, host, and port.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Dynamic [inbound | outbound] endpoint SSL configurations**. Then click the **New** button.

When an outbound connection is attempted, this association is checked ahead of the Secure Sockets Layer (SSL) configuration scope association. This means based on the target protocol, host, port, the outbound SSL configuration used can be different than the default specified in the SSL scope configuration.

Name:

Specifies the unique name of the dynamic endpoint configuration.

Data type: Text

Description:

Specifies text that describes the purpose of this dynamic selection criteria.

Data type: Text

Add connection information:

Specifies select information in the form protocol, host, port for the outbound connection. Multiple selection criteria can be entered. An asterisk (*) can be used to mean all protocols, hosts, or ports. You can use an * for any field.

Data type: Text

An example of selection criteria is `*,www.ibm.com,*`, which means that any time the target host is `www.ibm.com`, you must use the SSL configuration specified here. Another example selection criteria is `IIOP,*,*`, which means that any outbound IIOP request uses the SSL configuration that is specified in the SSL configuration field. When there is a conflict between two selection criteria, the application server uses the first match. The list of valid protocols you can use include: IIOP, HTTP, JMS, LDAP, SIP, ADMIN_SOAP, ADMIN_IIOP, or WEBSERVICES_HTTP.

Add:

Specifies to add the selected information from the **Add select information** menu to the right-hand list.

Remove:

Specifies to remove the selection from the right-hand list.

SSL Configuration:

Specifies the SSL configuration to be used by requests at this scope when a match occurs for the given selection criteria.

Data type: Text

Get certificate alias:

When selected, the keystore within the selected SSL configuration is queried for a list of personal certificates from which to choose.

Certificate alias:

Specifies the certificate alias that is used as the identity for the connection.

If you select **None**, the Java Secure Sockets Extension (JSSE) key manager determines which certificate is used. If multiple certificates exist in the keystore, the key manager might not consistently select the same certificate.

Data type: Text
Default: (none)

Quality of protection (QoP) settings

Use this page to specify security level, ciphers, and mutual authentication settings for the Secure Socket Layer (SSL) configuration.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **SSL configurations >** . Click on `{SSL_configuration_name}`. Under **Additional Properties**, click **Quality of protection (QoP) settings**.

Client authentication

Specifies the whether SSL client authentication should be requested if the SSL connection is used for the server side of the connection.

If None is selected, the server does not request that a client certificate be sent during the handshake. If Supported is selected, the server requests that a client certificate be sent. If the client does not have a certificate, the handshake might still succeed. If Required is selected, the server requests that a client certificate be sent. If the client does not have a certificate, the handshake fails.

Data type: Text
Default: None

Protocol

Specifies the Secure Sockets Layer (SSL) handshake protocol. This protocol is typically SSL_TLS, which supports all handshake protocols except for SSLv2 on the server side. When United States Federal Information Processing standard (FIPS) option is enabled, Transport Layer Security (TLS) is automatically used regardless of this setting.

Data type: text
Default: SSL_TLS

Predefined JSSE provider

Specifies one of the predefined Java Secure Sockets Extension (JSSE) providers. The IBMJSSE2 provider is recommended for use on all platforms which support it. It is required for use by the channel framework SSL channel. When Federal Information Processing Standard (FIPS) is enabled, IBMJSSE2 is used in combination with the IBMJCEPIS crypto provider.

Default: Enabled

Select provider

Specifies a package that implements a subset of the cryptography aspects for the Java security application programming interface (API). This value is a JSSE provider name that is listed in the java.security file. Note that cipher suites and protocol values depend upon the provider.

Data type: Text
Default: IBMJSSE2

Custom JSSE provider

Specifies that a custom JSSE provider should be used.

Default: Disabled

Custom provider

Specifies a package that implements a subset of the cryptography aspects for the Java security application programming interface (API). This value is a Java Secure Sockets Extension (JSSE) provider name that is listed in the java.security file. Note that cipher suites and protocol values depend upon the provider.

Data type: Text

Cipher suite groups

Specifies the various cipher suite groups that can be chosen depending upon your security needs. The stronger the cipher suite strength, the better the security; however, this can result in performance consequences.

Data type: Text
Default: Strong

Update selected ciphers

When selected, the cipher suites that are contained within the selected Cipher suite group are added to the list of **Selected ciphers**. Any change to this list changes the Cipher suite group to custom.

Selected ciphers

Specifies the ciphers that are effective when the configuration is saved. These ciphers are used to negotiate with the remote side of the connection during the handshake. A common cipher needs to be selected or the handshake fails.

Data type: Text

Add

Specifies to add the selected cipher to the **Selected ciphers** list.

Remove

Specifies to remove the selected cipher from the **Selected ciphers** list.

ssl.client.props client configuration file

Use the `ssl.client.props` file to configure Secure Sockets Layer (SSL) for clients. In previous releases of WebSphere Application Server, SSL properties were specified in the `sas.client.props` or `soap.client.props` files or as system properties. By consolidating the configurations, WebSphere Application Server enables you to manage security in a manner that is comparable to server-side configuration management. You can configure the `ssl.client.props` file with multiple SSL configurations.

Setting up the SSL configuration for clients

Client runtimes are dependent on the WebSphere Application Server `ssl.client.props` configurations.

Use the `setupCmdLine.bat` script on the command line to specify the `com.ibm.SSL.ConfigURL` system property.

```
-Dcom.ibm.SSL.ConfigURL=file:profile_root
\properties\ssl.client.props
```

The `com.ibm.SSL.ConfigURL` property references a file URL that points to the `ssl.client.props` file. You can reference the `CLIENTSSL` variable on the command line of any script that uses the `setupCmdLine.bat` file.

When you specify the `com.ibm.SSL.ConfigURL` system property, the SSL configuration is available to all protocols that use SSL. SSL configurations, which are referenced in the `ssl.client.props` file, also have aliases that you can reference. In the following sample code from the `sas.client.props` file, all of the SSL properties are replaced with a property that points to an SSL configuration in the `ssl.client.props` file:

```
com.ibm.ssl.alias=DefaultSSLSettings
```

The following sample code shows a property in the `soap.client.props` file that is similar to the `com.ibm.SSL.ConfigURL` property. This property references a different SSL configuration on the client side:

```
com.ibm.ssl.alias=DefaultSSLSettings
```

In the `ssl.client.props` file, you can change the administrative SSL configuration to avoid modifying the `soap.client.props` file.

Note: Support for SSL properties is still specified in the `sas.client.props` and `soap.client.props` files. However, consider moving the SSL configurations to the `ssl.client.props` file, because this file is the new configuration model for client SSL.

Properties of the ssl.client.props file

This section describes the default `ssl.client.props` file properties in detail, by sections within the file.

Global properties

Global SSL properties are process-specific properties that include Federal Information Processing Standard (FIPS) enablement, the default SSL alias, the `user.root` property for specifying the root location of the key and truststore paths, and so on.

Property	Default	Description
<code>com.ibm.ssl.defaultAlias</code>	<code>DefaultSSLSettings</code>	Specifies the default alias that is used whenever an alias is not specified by the protocol that calls the JSSEHelper API to retrieve an SSL configuration. This property is the final arbiter on the client side for determining which SSL configuration to use.

Property	Default	Description
com.ibm.ssl.validationEnabled	false	When set to true, this property validates each SSL configuration as it is loaded. Use this property for debug purposes only, to avoid unnecessary performance overhead during production.
com.ibm.ssl.performURLHostNameVerification	false	When set to true, this property enforces URL host name verification. When HTTP URL connections are made to target servers, the common name (CN) from the server certificate must match the target host name. Without a match, the host name verifier rejects the connection. The default value of false omits this check. As a global property, it sets the default host name verifier. Any javax.net.ssl.HttpURLConnection object can choose to enable host name verification for that specific instance by calling the setHostnameVerifier method with its own HostnameVerifier instance.
com.ibm.security.useFIPS	false	When set to true, FIPS-compliant algorithms are used for SSL and other Java Cryptography Extension (JCE)-specific applications. This property is typically not enabled unless the property is required by the operating environment.
user.root	C:\WebSphere\AppServer\profiles\default	This property can be used by key and truststore location properties as a single property for specifying the root path to the key and truststores. Typically, this property is the profile root. However, you can modify this property to any root directory on the local machine that has the proper read and potentially write authority to that directory.
profile_root of the profile	C:\WebSphere\AppServer\profiles\default	Key and truststore location properties can use this property as a single property for specifying the root path to the key and truststores. You can modify this property to any root directory on the local machine that has the proper read (and potentially write) authority to that directory.

Certificate creation properties

Use certificate creation properties to specify the default self-signed certificate values for the major attributes of a certificate. You can define the distinguished name (DN), expiration date, key size, and alias that are stored in the keystore.

Property	Default	Description
com.ibm.ssl.defaultCertReqAlias	default_alias	This property specifies the default alias to use to reference the self-signed certificate that is created in the keystore. If the alias already exists with that name, the default alias is appended with <code>_#</code> , where the number sign (#) is an integer that starts with 1 and increments until it finds a unique alias.
com.ibm.ssl.defaultCertReqSubjectDN	cn=\${hostname}, o=IBM,c=US	This property uses the property distinguished name (DN) that you set for the certificate when it is created. The <code>\${hostname}</code> variable is expanded to the host name on which it resides. You can use correctly formed DNs as specified by the X.509 certificate.
com.ibm.ssl.defaultCertReqDays	365	This property specifies the validity period for the certificate and can be as small as 1 day and as large as the maximum number of days that a certificate can be set, which is approximately 20 years.

Property	Default	Description
com.ibm.ssl.defaultCertReqKeySize	1024	This property is the default key size. The valid values depend upon the Java Virtual Machine (JVM) security policy files that are installed. By default, the product JVMs ship with the export policy file that limits the key size to 1024. To get a large key size such as 2048, you can download the restricted policy files from the Web site.

Certificate revocation checking

To enable certificate revocation checking, you can set a combination of Online Certificate Status Protocol (OCSP) properties. These properties are not used unless you set the `com.ibm.ssl.trustManager` property to `IbmPKIX`. In addition, to successfully process revocation checking on the client, you must turn off the signer exchange prompt. To turn off the signer exchange prompt, change the `com.ibm.ssl.enableSignerExchangePrompt` property to `false`. For more information, see the related link to the "Enabling certificate revocation checking with the default `IbmPKIX` trust manager" topic.

SSL configuration properties

Use the SSL configuration properties section to set multiple SSL configurations. For a new SSL configuration specification, set the `com.ibm.ssl.alias` property because the parser starts a new SSL configuration with this alias name. The SSL configuration is referenced by using the alias property from another file, such as `sas.client.props` or `soap.client.props`, through the default alias property. The properties that are specified in the following table enable you to create a `javax.net.ssl.SSLContext`, among other SSL objects.

Property	Default	Description
com.ibm.ssl.alias	DefaultSSLSettings	This property is the name of this SSL configuration and must be the first property for an SSL configuration because it references the SSL configuration. If you change the name of this property after it is referenced elsewhere in the configuration, the runtime defaults to the <code>com.ibm.ssl.defaultAlias</code> property whenever the reference is not found. The error <code>trust file is null</code> or <code>key file is null</code> might display when you start an application using an SSL reference that is no longer valid.
com.ibm.ssl.protocol	SSL_TLS	This property is the SSL handshake protocol that is used for this SSL configuration. This property attempts Transport Layer Security (TLS) first, but accepts any remote handshake protocol, including SSLv3 and TLS. Valid values for this property include SSLv2 (client side only), SSLv3, SSL, TLS, TLSv1, and SSL_TLS.
com.ibm.ssl.securityLevel	HIGH	This property specifies the cipher group that is used for the SSL handshake. The typical selection is HIGH, which specifies 128-bit or higher ciphers. The MEDIUM selection provides 40-bit ciphers. The LOW selection provides ciphers that do not perform encryption, but do perform signing for data integrity. If you specify your own cipher list selection, uncomment the property <code>com.ibm.ssl.enabledCipherSuites</code> .

Property	Default	Description
com.ibm.ssl.trustManager	lbmX509	This property specifies the default trust manager that you must use to validate the certificate sent by the target server. This trust manager does not perform certificate revocation list (CRL) checking. You can choose to change this value to lbmPKIX for CRL checking using CRL distribution lists in the certificate, which is a standard way to perform CRL checking. When you want to perform custom CRL checking, you must implement a custom trust manager and specify the trust manager in the com.ibm.ssl.customTrustManagers property. The lbmPKIX option might affect performance because this option requires IBM CertPath for trust validation. Use lbmX509 unless CRL checking is necessary. If you are using the Online Certificate Status Protocol (OCSP) properties, set this property value to lbmPKIX.
com.ibm.ssl.keyManager	lbmX509	This property specifies the default key manager to use for choosing the client alias from the specified keystore. This key manager uses the com.ibm.ssl.keyStoreClientAlias property to specify the keystore alias. If this property is not specified, the choice is delegated to Java Secure Socket Extension (JSSE). JSSE typically chooses the first alias that it finds.
com.ibm.ssl.contextProvider	IBMJSSE2	This property is used to choose the JSSE provider for the SSL context creation. It is recommended that you default to IBMJSSE2 when you use a Java virtual machine (JVM). The client plug-in can use the SunJSSE provider when using a Sun JVM.
com.ibm.ssl.enableSignerExchangePrompt	true	This property determines whether to display the signer exchange prompt when a signer is not present in the client truststore. The prompt displays information about the remote certificate so that WebSphere Application Server can decide whether or not to trust the signer. It is very important to validate the certificate signature. This signature is the only reliable information that can guarantee that the certificate has not been modified from the original server certificate. For automated scenarios, disable this property to avoid SSL handshake exceptions. Run the retrieveSigners script, which sets up the SSL signer exchange, to download the signers from the server prior to running the client. If you are using the Online Certificate Status Protocol (OCSP) properties, set this property value to false.
com.ibm.ssl.keyStoreClientAlias	default	This property is used to reference an alias from the specified keystore when the target does not request client authentication. When WebSphere Application Server creates a self-signed certificate for the SSL configuration, this property determines the alias and overrides the global com.ibm.ssl.defaultCertReqAlias property.
com.ibm.ssl.customTrustManagers	Commented out by default	This property enables you to specify one or more custom trust managers, which are separated by commas. These trust managers can be in the form of <i>algorithm provider</i> or <i>classname</i> . For example, lbmX509 IBMJSSE2 is in the <i>algorithm provider</i> format, and the com.acme.myCustomTrustManager interface is in the <i>classname</i> format. The class must implement the javax.net.ssl.X509TrustManager interface. Optionally, the class can implement the com.ibm.wsspi.ssl.TrustManagerExtendedInfo interface. These trust managers run in addition to the default trust manager that is specified by the com.ibm.ssl.trustManager interface. These trust managers do not replace the default trust manager.

Property	Default	Description
com.ibm.ssl.customKeyManager	Commented out by default	This property enables you to have one, and only one, custom key manager. The key manager replaces the default key manager that is specified in the com.ibm.ssl.keyManager property. The form of the key manager is <i>algorithm provider</i> or <i>classname</i> . See the format examples for the com.ibm.ssl.customTrustManagers property. The class must implement the javax.net.ssl.X509KeyManager interface. Optionally, the class can implement the com.ibm.wsspi.ssl.KeyManagerExtendedInfo interface. This key manager is responsible for alias selection.
com.ibm.ssl.dynamicSelectionInfo	Commented out by default	This property enables dynamic association with the SSL configuration. The syntax for a dynamic association is <i>outbound_protocol, target_host, or target_port</i> . For multiple specifications, use the vertical bar () as the delimiter. You can replace any of these values with an asterisk (*) to indicate a wildcard value. Valid <i>outbound_protocol</i> values include: IIOP, HTTP, LDAP, SIP, BUS_CLIENT, BUS_TO_WEBSPHERE_MQ, BUS_TO_BUS, and ADMIN_SOAP. When you want the dynamic selection criteria to choose the SSL configuration, uncomment the default property, and add the connection information. For example, add the following on one line com.ibm.ssl.dynamicSelectionInfo=HTTP, .ibm.com,443 HTTP,.ibm.com,9443
com.ibm.ssl.enabledCipherSuites	Commented out by default	This property enables you to specify a custom cipher suite list and override the group selection in the com.ibm.ssl.securityLevel property. The valid list of ciphers varies according to the provider and JVM policy files that are applied. For cipher suites, use a space as the delimiter.
com.ibm.ssl.keyStoreName	ClientDefaultKeyStore	This property references a keystore configuration name. If you have not already defined the keystore, the rest of the keystore properties must follow this property. After you define the keystore, you can specify this property to reference the previously specified keystore configuration. New keystore configurations in the ssl.client.props file have a unique name.
com.ibm.ssl.trustStoreName	ClientDefaultTrustStore	This property references a truststore configuration name. If you have not already defined the truststore, the rest of the truststore properties must follow this property. After you define the truststore, you can specify this property to reference the previously specified truststore configuration. New truststore configurations in the ssl.client.props file should have a unique name.

Keystore configurations

SSL configurations reference keystore configurations whose purpose is to identify the location of certificates. Certificates represent the identity of clients that use the SSL configuration. You can specify keystore configurations with other SSL configuration properties. However, it is recommended that you specify the keystore configurations in this section of the ssl.client.props file after the com.ibm.ssl.keyStoreName property identifies the start of a new keystore configuration. After you fully define the keystore configuration, the com.ibm.ssl.keyStoreName property can reference the keystore configuration at any other point in the file.

Property	Default	Description
com.ibm.ssl.keyStoreName	ClientDefaultKeyStore	This property specifies the name of the keystore as it is referenced by the runtime. Other SSL configurations can reference this name further down in the ssl.client.props file to avoid duplication.

Property	Default	Description
com.ibm.ssl.keyStore	\${user.root}/etc/ key.p12	This property specifies the location of the keystore in the required format of the com.ibm.ssl.keyStoreType property. Typically, this property references a keystore file name. However, for cryptographic token types, this property references a Dynamic Link Library (DLL) file.
com.ibm.ssl.keyStorePassword	WebAS	This property is the default password, which is the cell name for the profile when it is created. The password is typically encoded using an {xor} algorithm. You can use iKeyman to change the password in the keystore, then change this reference. If you do not know the password and if the certificate is created for you, change the password in this property, then delete the keystore from the location where it resides. Restart the client to recreate the keystore by using the new password, but only if the keystore name ends with DefaultKeyStore and if the fileBased property is true. Delete both the keystore and truststore at the same time so that a proper signer exchange can occur when both are recreated together.
com.ibm.ssl.keyStoreType	PKCS12	This property is the keystore type. Use the default, PKCS12, because of its interoperability with other applications. You can specify this property as any valid keystore type that is supported by the JVM on the provider list.
com.ibm.ssl.keyStoreProvider	IBMJCE	The IBM Java Cryptography Extension property is the keystore provider for the keystore type. The provider is typically IBMJCE or IBMPKCS11Impl for cryptographic devices.
com.ibm.ssl.keyStoreFileBased	true	This property indicates to the runtime that the keystore is file-based, meaning it is located on the file system.
com.ibm.ssl.keyStoreReadOnly	false	This property indicates to the run time for WebSphere Application Server whether the key store can be modified during the run time.

Truststore Configurations

SSL configurations reference truststore configurations, whose purpose is to contain the signer certificates for servers that are trusted by this client. You can specify these properties with other SSL configuration properties. However, it is recommended that you specify truststore configurations in this section of the `ssl.client.props` file after the `com.ibm.ssl.trustStoreName` property has identified the start of a new truststore configuration. After you fully define the truststore configuration, the `com.ibm.ssl.trustStoreName` property can reference the configuration at any other point in the file.

A truststore is a keystore that JSSE uses for trust evaluation. A truststore contains the signers that WebSphere Application Server requires before it can trust the remote connection during the handshake. If you configure the `com.ibm.ssl.trustStoreName=ClientDefaultKeyStore` property, you can reference the keystore as the truststore. Further configuration is not required for the truststore because all of the signers that are generated through signer exchanges are imported into the keystore where they are called by the runtime.

Property	Default	Description
com.ibm.ssl.trustStoreName	ClientDefaultTrustStore	This property specifies the name of the truststore as it is referenced by the runtime. Other SSL configurations can reference further down in the <code>ssl.client.props</code> file to avoid duplication.
com.ibm.ssl.trustStore	<code>\${user.root}/etc/trust.p12</code>	This property specifies the location of the truststore in the format that is required by the truststore type that is referenced by the <code>com.ibm.ssl.trustStoreType</code> property. Typically, this property references a truststore file name. However, for cryptographic token types, this property references a DLL file.
com.ibm.ssl.trustStorePassword	WebAS	This property specifies the default password, which is the cell name for the profile when it is created. The password is typically encoded using an {xor} algorithm. You can use iKeyman to change the password in the keystore, then change the reference in this property. If you do not know the password and if the certificate was created for you, change the password in this property, then delete the truststore from the location where it resides. Restart the client to recreate the truststore by using the new password, but only if the keystore name ends with <code>DefaultTrustStore</code> and the <code>fileBased</code> property is <code>true</code> . It is recommended that you delete the keystore and the truststore at the same time so that a proper signer exchange can occur when both are recreated together.
com.ibm.ssl.trustStoreType	PKCS12	This property is the truststore type. Use the default PKCS12 type because of its interoperability with other applications. You can specify this property as any valid truststore type that is supported by the JVM functionality on the provider list.
com.ibm.ssl.trustStoreProvider	IBMJCE	This property is the truststore provider for the truststore type. The provider is typically <code>IBMJCE</code> or <code>IBMPKCS11Impl</code> for cryptographic devices.
com.ibm.ssl.trustStoreFileBased	true	This property indicates to the runtime that the truststore is file-based, meaning it is located on the file system.
com.ibm.ssl.trustStoreReadOnly	false	This property indicates to the run time for WebSphere Application Server whether the truststore can be modified during the run time.

Create a CA client

A plug point is provided to allow users to connect to a certificate authority (CA) to request, query, and revoke certificates. A security configuration object, called a CAClient, must be created for WebSphere to communicate with the CA. The CAClient object must contain a WSPKIClient() implementation, and it will handle the connection and communicate with the CA server. Users can also create their own implementation.

Before you begin

The WSPKIClient interface must be implemented and the class name provided as part of the CAClient when it is created.

You use the administrative console to create a new CA client.

1. Click **Security > SSL certificate and key management**.
2. Click **Certificate Authority (CA) client configurations**. A panel of existing CA clients appears.
3. Click **New** to create a new CA client in the configuration.

Note: You can also create a CA client by using the **createCAClient** AdminTask .

4. Fill in the following information for the CA client
 - Name of the CA client.
 - The management scope (selected from the drop-down list).
 - WSPKIClient implementation class.
 - CA server host name.
 - User name.
 - Password.
 - Confirm of password.
 - Number of times to poll.
 - Polling interval (in minutes) when requesting certificates.
 - Custom properties.
5. Click **Apply** then **OK**.

Results

The information in the object can then be used by the runtime to connect to a CA to create, revoke, or replace a certificate.

Delete a CA client

You can delete the CAClient object from the security configuration if a connection to a certificate authority (CA) is no longer needed.

Before you begin

You use the administrative console to delete a CA client.

1. Click **Security > SSL certificate and key management**.
2. Click **Certificate Authority (CA) client configurations**. A panel displaying the existing CA clients appears.
3. Click the CA client name you want to delete.
4. Click the **Delete** button.

Note: You can also use the **deleteCAClient** AdminTask to delete the CA client.

Results

The CA client is deleted from the configuration.

Note: When you use the **deleteCAClient** AdminTask to delete the CA client, the CA client cannot be deleted if a CA certificate that exists in the keystore was obtained from the certificate authority and is still referenced by the CA client. For example, when such CA certificate still exists, the user receives the following message:

```
wsadmin>$AdminTask deleteCAClient {-caClientName myca}
WASX7015E: Exception running command:
"$AdminTask deleteCAClient {-caClientName myca}"; exception information:
com.ibm.websphere.management.cmdframework.CommandValidationException:
CWPKI0687E: The Certificate Authority (CA) client myca is still referenced by:
[Certificate alias myca21 in key store CellDefaultKeyStore].
wsadmin>
```

View or modify a CA client

You can view or modify the CAClient object settings in the security configuration. The CAClient object contains all the information needed to connect and communicate with a certificate authority (CA). A connection to a Certificate Authority is used to request a certificate, query a certificate, or revoke a certificate.

Before you begin

You use the administrative console to view or modify a CA client.

1. Click **Security > SSL certificate and key management**.
2. Click **Certificate Authority (CA) client configurations**. A panel displaying the existing CA clients appears.
3. Click the CA client name you want to examine and modify.

Note: You can also use the **getCAClient** AdminTask to get information about the existing CA client and the **modifyCAClient** AdminTask to make changes to the CA client.

4. Make the changes to the CA client information as required. Modify the following information as required.
 - Name of the CA client.
 - The management scope (selected from the drop-down list).
 - Implementation class.
 - CA server host name.
 - User name.
 - Password.
 - Confirm of password.
 - Number of times to poll.
 - Polling interval (in minutes) when requestin certificates.
 - Custom properties.
5. Click **Apply** then **OK**.

Results

The information in the object can then be used by the runtime to connect to a CA to create, revoke, or replace a certificate

What to do next

Creating a keystore configuration for a preexisting keystore file

A Secure Sockets Layer (SSL) configuration references keystore configurations during WebSphere Application Server runtime. Whether a keystore file was created by another keystore tool or saved from a previous configuration, the file must be referenced by a keystore configuration object to be used by the server. A keystore configuration object can be created to reference a pre-existing keystore object.

Before you begin

A keystore must already exist.

Note: To create a keystore by using the wsadmin tool, use the **createKeyStore** command of the AdminTask object. For more information, see “KeyStoreCommands command group for the AdminTask object” on page 981

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound}**.
2. Under Related Items, click **Key stores and certificates**, then click **New**.
3. Type a name in the **Name** field. This name uniquely identifies the keystore in the configuration.
4. Type the location of the keystore file in the **Path** field. The location can be a file name or a file URL to an existing keystore file.
5. Type the keystore password in the **Password** field. This password is for the keystore file that you specified in the **Path** field.
6. Type the keystore password again in the **Confirm Password** field to confirm the password.
7. Select a keystore type from the list. The type that you select is for the keystore file that you specified in the **Path** field.
8. Select any of the following optional selections:
 - The **Read only selection** creates a keystore configuration object but does not create a keystore file. If this option is selected, the keystore file that you specified in the **Path** field must already exist.
 - The **Initialize at startup selection** initializes the keystore during runtime.
 - The **Enable cryptographic operations on a hardware device** specifies whether a hardware cryptographic device is used for cryptographic operations only.

Note: Operations that require login are not supported when using this option.

9. Click **Apply** and **Save**.

Results

You have created a keystore configuration object for the keystore file that you specified. This keystore can now be used in an SSL configuration.

What to do next

You can create additional keystore configurations, as needed.

Configuring a hardware cryptographic keystore

You can create a hardware cryptographic keystore that WebSphere Application Server can use to provide cryptographic token support in the server configuration.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Key stores and certificates**.
2. Click **New**.
3. Type a name to identify the keystore. This name is used to enable hardware cryptography in the Web services security configuration.
4. Optionally, you can type a description for the keystore in the **Description** field.
5. You can specify a **Management scope** for the key store. This is not required. The management scope specifies the scope where this Secure Sockets Layer (SSL) configuration is visible. For example, if you choose a specific node, then the configuration is only visible on that node and any servers that are part of that node.
6. Type the path for the hardware device-specific configuration file. The configuration file is a text file that contains entries in the following format: *attribute = value*. The valid values for attribute and value are described in detail in the Software Developer Kit, Java Technology Edition documentation. The two mandatory attributes are name and library, as shown in the following sample code:

```
name = FooAccelerator
library = /opt/foo/lib/libpkcs11.so
slotListIndex = 0
```

The configuration file should also include device-specific configuration data. Navigate to the PKCS11ImplConfigSamples.jar file, which contains sample configuration files, under the heading "PKCS 11 Implementation Provider" on the Java technology site <http://www.ibm.com/developerworks/java/jdk/security/50/>.

Note: If you want to use the IBMPKCS11Impl provider, you must initialize the provider individually and explicitly express the provider in the JCE getInstance method. JSEE2 is unable to use the IBMPKCS11Impl provider for acceleration.

- a. You can use this link <http://www.ibm.com/developerworks/java/jdk/security/50/secguides/pkcs11implDocs/IBMJavaPKCS11ImplementationProvider.html> to initialize the IBMPKCS11 provider in a thread safe way
- b. Specify a unique .cfg file that contains information about the supported hardware device. A list of supported hardware devices are available at <http://www.ibm.com/developerworks/java/jdk/security/50/secguides/pkcs11implDocs/IBMPKCS11SupportList.html>
- c. You specify the Signature.getInstance method with the properly initialized IBMPKCS11Impl provider instance as shown.

```
Signature.getInstance("SHA1withRSA", ibmpkcs11implinstance);
```

7. Type a password if the token login is required. Operations that use keys on the token require a secure login. This field is optional if the keystore is used as a cryptographic accelerator. In this case, you need to select **Enable cryptographic operations on hardware device**.
8. Select the **PKCS11** type.
9. Select **Read only**.
10. Click **OK** and **Save**.

Results

WebSphere Application Server can now provide cryptographic token support in the server configuration.

What to do next

You can refer to this keystore in any server Secure Sockets Layer (SSL) configuration to achieve the following results:

- Cryptographic acceleration because the cryptographic hardware device has no persistent key storage
- Secure cryptographic hardware because a cryptographic token generates and securely stores the private key that WebSphere Application Server uses for SSL key exchange.

You can also refer to this keystore in the Web services security default bindings configuration to achieve similar results.

Managing keystore configurations remotely

You can manage keystores remotely in a Network Deployment environment on separate machines. A node server can hold the configuration for a keystore, while the actual keystore resides on another system. After you set up a remotely managed configuration, you can perform all of the certificate and keystore operations for the keystore on the remote machine from the server that contains the keystore remote configuration.

Before you begin

Key stores can be remotely managed only in network deployed environments.

Note: To manage a self-signed certificates by using the wsadmin tool, use the **PersonalCertificateCommands** group commands of the AdminTask object. For more information, see “PersonalCertificateCommands command group for the AdminTask object” on page 1022.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates**.
2. Click **New**.
3. Type a name in the **Name** field. This name uniquely identifies the keystore in the configuration.
4. Type the location of the keystore file in the **Path** field. The location can be a file name or a file Uniform Resource Locator (URL) to an existing keystore file.
5. Type the keystore password in the **Password** field. This password is for the keystore file that you specified in the **Path** field.
6. Type the keystore password again in the **Confirm Password** field to confirm the password.
7. Select a keystore type from the list. The type you select is for the keystore file that you specified in the **Path** field.
8. Select the **Remotely managed** check box, and then fill in one or more hosts names of the systems where the keystore file is to be located. If you provide multiple host names, separate the host names with a pipe (|).
9. Select any of the following optional selections:
 - The **Read only selection** creates a keystore configuration object but does not create a keystore file. If this option is selected, the keystore file that you specified in the **Path** field must already exist.
 - The **Initialize at startup selection** initializes the keystore during run time.
10. Select **Apply** and **Save**.

Results

A keystore configuration object is created on the server from where the command was run. The keystore file for the configuration will be created on each system that you specified in the host list.

What to do next

Now, you can perform all certificate management operations on the keystore from the system where the keystore configuration resides. For example, you can perform certificate management operations, such as: creating a self-signed certificate, extracting a certificate, or extracting a signer certificate.

Keystores and certificates collection

Use this page to manage keystore types, including cryptography, Resource Access Control Facility (RACF), Certificate Management Services (CMS), Java, and all trust store types.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > *ssl_configuration***. Under Related items, click **Keystores and certificates**.

Button	Resulting action
New	Adds a new keystore object that can be referenced by Secure Sockets Layer (SSL) configurations or KeySets. The Keystore management scope is based on the part of the topology tree from which it was created.
Delete	Deletes an existing keystore. The keystore should not be referenced by any other parts of the configuration before you delete it.
Change password	Allows for changing a keystore password.
Exchange signers	Refers to exchanging signers in a keystore. You can select two keystores, along with personal certificates or signer certificates from a selected keystore, then add them as a signer to another selected keystore.

Keystore usages

Filters the keystore usage types in the keystore collection.

The default value for the keystore usage filter depends on the navigation path that you followed to get to the Keystores and certificates panel. You can change the value of the keystore usage filter by clicking on the drop-down list and selecting a different filter value.

Navigation path	Keystore usage default value
Security > Global security > Administrative authentication > Manage administrative keystore > Keystores and certificates	RSA token keystores
Security > SSL certificate and key management > Keystores and certificates	SSL keystores
Security > SSL certificate and key management > Key sets > CellLTPAKeyPair > Keystores and certificates	Key set keystores
Security > SSL certificate and key management > SSL configurations > CellDefaultSSLSettings > Keystores and certificates	SSL keystores
Security > SSL certificate and key management > Manage endpoint security configurations > <i>node name</i> > Keystores and certificates	SSL keystores

Name

Specifies the unique name that is used to identify the keystore. This name is typically scoped by the ManagementScope scopeName and based upon the location of the keystore. The name must be unique within the existing keystore collection.

This is a user-defined name.

Description

Specifies the description of the keystore.

This is a user-defined description.

Path

Specifies the location of the keystore file in the format needed by the keystore type. This file can be a card-specific configuration file for cryptographic devices or a filename or file URL for file-based keystores. It can be a safkeyring URL for RACF keyrings.

Key store settings

Use this page to create all keystore types, including cryptographic, Resource Access Control Facility (RACF), Certificate Management Services (CMS), Java, and all truststore types.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound}**. Under Related Items, click **Key stores and certificates**. Click either **New** or an existing keystore.

Links to Personal certificates, Signer certificates, and Personal certificate requests enable you to manage certificates in a manner similar to iKeyman capabilities. A keystore can be file-based, such as CMS or Java keystore types, or it can be remotely managed.

Note: Any changes made to this panel are permanent.

Name

Specifies the unique name to identify the keystore. The keystore is typically scoped by the ManagementScope scopeName based on the location of the keystore. The name must be unique within the existing keystore collection.

Data type: Text

Description

Specifies the description of the keystore.

Data type: Text

Management scope

Specifies the scope where this Secure Sockets Layer (SSL) configuration is visible. For example, if you choose a specific node, then the configuration is only visible on that node and any servers that are part of that node.

Data type: Text

Path

Specifies the location of the keystore file in the format needed by the keystore type. This file can be a dynamic link library (DLL) for cryptographic devices or a filename or file URL for file-based keystores. It can be a safkeyring URL for RACF keyrings.

Data type: Text

Control region user

Specifies the Control region Started Task user ID in which the Control region System Authorization Facility (SAF) keyring is created. The user ID must match the exact ID being used by the Control region. Note: This option only applies when creating writable SAF keyrings on z/OS.

Data type: Text

Servant region user

Specifies the Servant region Started Task user ID in which the Servant region System Authorization Facility (SAF) keyring is created. The user ID must match the exact ID being used by the Servant region. Note: This option only applies when creating writable SAF keyrings on z/OS.

Data type: Text

Password [new keystore] | Change password [existing keystore]

Specifies the password used to protect the keystore. For the default keystore (names ending in DefaultKeyStore or DefaultTrustStore), the password is WebAS. This default password must be changed.

This field can be edited.

Data type: Text

Confirm password

Specifies confirmation of the password to open the keystore file or device.

Data type: Text

Type

Specifies the implementation for keystore management. This value defines the tool that operates on this keystore type.

The list of options is returned by `java.security.Security.getAlgorithms("KeyStore")`. Some options might be filtered and some might be added based on the `java.security` configuration.

Data type: Text
Default: PKCS12

Read only

Specifies whether the keystore can be written to or not. If the keystore cannot be written to, certain operations cannot be performed, such as creating or importing certificates.

Default: Disabled

Remotely managed

Specifies whether the key store is remotely managed, which means that a remote MBean call is needed to update the key store based on the host name specified in the host list field. Most hardware cryptographic token devices are remotely managed. If a key store is marked remotely managed, list the host name of the server where the device is installed in the Host list field.

Default:

Initialize at startup

Specifies whether the keystore needs to be initialized before it can be used for cryptographic operations. If enabled, the keystore is initialized at server startup.

Default: Disabled

Enable cryptographic operations on hardware device

Specifies whether a hardware cryptographic device is used for cryptographic operations only. Operations that require a login are not supported when using this option.

Default: Disabled

Key managers collection

Use this page to define the implementation settings for key managers. A key manager is invoked during a Secure Sockets Layer (SSL) handshake to determine which certificate alias is used. The default key manager (WSX509KeyManager) performs alias selection. If more advanced function is desired, define a custom key manager on the **Manage endpoint security configurations** panel.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key managers**.

Button	Resulting action
New	Adds a new key manager that can be selected by an SSL configuration. A key manager is invoked during an SSL handshake to select a specific certificate alias to use from a key store.
Delete	Deletes an existing key manager. The key manager should not be referenced by any SSL configuration before you can delete it.

Name

Specifies the name of the key manager, which you can select on the SSL configuration panel.

Class name

Specifies the name of the key manager implementation class. This class implements javax.net.ssl.X509KeyManager interface and, optionally, the com.ibm.wsspi.ssl.KeyManagerExtendedInfo interface.

Algorithm

Specifies the algorithm name of the key manager that is implemented by the selected provider.

Key managers settings

Use this page to define key managers implementation settings. A key manager gets invoked during an Secure Sockets Layer (SSL) handshake to determine the certificate alias to be used. The default key manager (WSX509KeyManager) performs alias selection. If more advanced function is desired, a custom key manager can be specified here and selected in the SSL configuration.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key managers**. Then click the **New** button.

Name

Specifies the name of the key manager, which you can select on the SSL configuration panel.

Data type: Text

Standard

Specifies the key manager selection that is available from a Java provider that is installed in the java.security file. This provider might be shipped by Java Secure Sockets Extension (JSSE) or be a custom provider that implements an X509KeyManager interface.

Default: Enabled

Provider

Specifies the provider name that has an implementation of an X509KeyManager interface. This provider is typically set to IBMJSSE2.

Data type: Text

Default: IBMJCE

Algorithm

Specifies the algorithm name of the trust manager implemented by the selected provider.

Data type: Text

Default: lbmX509

Custom

Specifies that the key manager selection is based on a custom implementation class that implements the javax.net.ssl.X509KeyManager interface and optionally the com.ibm.wsspi.ssl.KeyManagerExtendedInfo interface to obtain additional connection information not otherwise available.

Default: Disabled

Class name

Specifies the name of the key manager implementation class.

Data type: Text

Creating a self-signed certificate

You can create a self-signed certificate. WebSphere Application Server uses the certificate at runtime during the handshake protocol. Self-signed certificates are located in the default keystore.

Before you begin

You must create a keystore before you can create a self-signed certificate.

Note: To create a self-signed certificate by using the wsadmin tool, use the **createSelfSignedCertificate** command of the AdminTask object. For more information, see “PersonalCertificateCommands command group for the AdminTask object” on page 1022.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates > [keystore]**.
2. From Additional Properties, click **Personal certificates**.
3. Click **Create a self-signed certificate**.
4. Type a certificate alias name. The alias identifies the certificate request in the keystore.
5. Type a common name (CN) value. This value is the CN value in the certificate distinguished name (DN).
6. Type the validity period The default validity period value is 365 days.
7. You can configure one or more of the following optional values:
 - a. Optional: Select a key size value. The default key size value is 1024 bits.
 - b. Optional: Type an organization value. This value is the O value in the certificate DN.
 - c. Optional: Type an organizational unit value. This organizational unit value is the OU value in the certificate DN.
 - d. Optional: Type a locality value. This locality value is the L value in the certificate DN.
 - e. Optional: Type a state or providence value. This value is the ST value in the certificate DN.
 - f. Optional: Type a zip code value. This zip code value is the POSTALCODE value in the certificate DN.
 - g. Optional: Select a country value from the list. This country value is the C= value in the certificate request DN.
8. Click **Apply**.

Results

You have created a self-signed certificate that resides in the keystore. The SSL configuration for the WebSphere Application Server runtime uses this certificate for SSL communication. Extract the signer of the self-signed certificate to add the signer to another keystore.

Replacing an existing self-signed certificate

Occasionally, you need to replace an existing or expired self-signed certificate with a new certificate. Certificates are referenced in the runtime configuration by the Secure Sockets Layer (SSL) Configuration object and the Dynamic SSL Configuration Selection object. You can replace a certificate with a new certificate alias reference or with a new signer certificate.

Before you begin

The current certificate and the certificate replacement must exist in the same keystore before you can replace a certificate.

Note: To replace a self-signed certificate by using the wsadmin tool, use the **replaceCertificate** command of the AdminTask object. For more information, see `rxml_atpersonalcert.dita`.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates > [keystore]**.
2. Under Additional Properties, click **Personal certificates**.
3. Select the certificate to be replaced. The alias list must include the certificate to be replaced and the certificate to replace it with.
4. Click **Replace**.
5. Select a replacement certificate alias from the list.

6. You can delete one of the following types of certificates:
 - Select **Delete old certificate** to delete the existing or expired certificate.
 - Select **Delete old signers** to delete the existing signer certificates.
7. Click **Apply**.

Results

Your results depend on what you selected:

- If you selected **Delete old certificate**, the new certificate alias replaces all of the references to the certificate alias in the configuration.
- If you selected **Delete old signers**, the new signer certificate replaces all of the occurrences of the old signer certificates.
- If the new certificate alias replaces the existing alias, the WebSphere Application Server runtime checks to make sure that:
 - All of the SSL Configurations objects reference the certificate
 - The Dynamic SSL Configuration Selections objects and the SSL Configuration group objects reference the certificate.
- If you selected **Delete old signers**, the existing signer certificates are replaced.
- If you selected **Delete old certificate**, the existing certificate is deleted.

Creating a certificate authority request

To ensure Secure Sockets Layer (SSL) communication, servers require a personal certificate that is either self-signed, chained or signed by an external certificate authority (CA). You must first create a personal certificate request to obtain a certificate that is signed by a CA.

Before you begin

The keystore that contains a personal certificate request must already exist.

Note: To create a certificate request by using the wsadmin tool, use the **createCertificateRequest** command of the AdminTask object. For more information, see “CertificateRequestCommands command group of the AdminTask object” on page 1049.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Key stores and certificates > keystore**.
2. Click **Personal certificate requests > New**.
3. Type the full path of the certificate request file. The certificate request is created in this location.
4. Type an alias name in the **Key label** field. The alias identifies the certificate request in the keystore.
5. Type a common name (CN) value. This value is the CN value in the certificate distinguished name (DN).
6. You can configure one or more of the following optional values:
 - a. Optional: Select a key size value. The default key size value is 1024 bits.
 - b. Optional: Type an organization value. This value is the O value in the certificate DN.
 - c. Optional: Type an organizational unit value. This organizational unit value is the OU value in the certificate DN.
 - d. Optional: Type a locality value. This locality value is the L value in the certificate DN.
 - e. Optional: Type a state or providence value. This value is the ST value in the certificate DN.

- f. Optional: Type a zip code value. The zip code value is the POSTALCODE value in the certificate DN.
 - g. Optional: Select a country value from the list. This country value is the C= value in the certificate request DN.
7. Click **Apply**.

Results

The certificate request is created in the specified file location in the keystore. The request functions as a temporary placeholder for the signed certificate until you manually receive the certificate in the keystore.

Note: Key store tools (such as iKeyman and keyTool) cannot receive signed certificates that are generated by certificate requests from WebSphere Application Server. Similarly, WebSphere Application Server cannot accept certificates that are generated by certificate requests from other keystore utilities.

What to do next

Now you can receive the CA-signed certificate into the keystore to complete the process of generating a signed certificate for your server.

Certificate request settings

Use this page to verify the properties of a personal certificate request.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificate requests > certificate request**.

File for certificate request

Specifies the fully qualified name of the file that contains the certificate request that can be sent to a certificate authority (CA) server.

Key label

Specifies the certificate alias name for the signer in the key store.

Key size

Specifies the size of the keys that are generated.

Requested by

Specifies the Subject distinguished name (DN) that represents the identity of the certificate request.

Fingerprint (SHA Digest)

Specifies the SHA hash of the personal certificate, which can be used to verify that the certificate has not been altered when it is used in a remote connection.

Signature algorithm

Specifies the algorithm used to sign the certificate.

Personal certificates collection

Use this page to manage personal certificates.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificates**.

The **Personal certificates** page lists all personal certificates in the selected key store. You can do most certificate management operations in this panel, including creating a new self-signed certificate, deleting a certificate, receiving one generated from a CA, replacing a certificate (simultaneous delete and create, replacing references across all key stores), extracting the signer, and importing or exporting a personal certificate.

Personal certificate requests are temporary place holders for certificates that will be signed by a certificate authority (CA).

The Key store collection must contain at least two key store files. You must select one file in order to replace, extract, or export a key store,

Button	Resulting action
Create (drop-down list)	Enables the application server to create the following certificates: <ul style="list-style-type: none"> • Self-signed Certificate • CA-signed Certificate • Chained Certificate
Delete	Specifies to delete a certificate from the key store. Be careful that the certificate alias is not referenced elsewhere in the Secure Sockets Layer configuration.
Receive a certificate from a certificate authority	Enables the application server to receive a certificate authority (CA)-generated certificate from a file to complete a certificate request.
Replace	Replaces a personal certificate with another personal certificate. All key stores in the configuration looking for signer certificate form the original personal certificate and replaces them with the new personal certificates signer. Any place in the security configuration where the certificate alias is referenced will be replaced with the new certificate alias.
Extract	Extracts the signer part of personal certificate from the key store and stores it to a file. The file can then be used to add the signer to another key store.
Import	Imports a certificate, including the private key, from a key store file or managed key store.
Export	Exports a certificate, including the private key, to a specified key store file or manage key store.
Revoke	Revokes a CA-signed certificate.
Renew	Renews a self signed or chained certificate.

Alias

Specifies the alias by which the personal certificate is referenced in the key store.

When you select an alias, the View Certificate panel opens.

Issued by

Specifies the distinguished name of the entity by which the certificate was issued. This name is the same as the issued-to distinguished name when the personal certificate is self-signed.

Issued to

Specifies the distinguished name of the entity to which the certificate was issued.

Serial number

Specifies the certificate serial number that is generated by the issuer of the certificate.

Expiration

Specifies the expiration date of the signer certificate for validation purposes.

Self-signed certificates settings

Use this page to create self-signed certificates.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > keystore**. Under Additional Properties, click **Personal certificates > Create (drop-down list) > Self-signed certificate**.

This same help file is available when you create a new certificate or view an existing certificate. The fields in this help file are described according to how they appear and are used on the administrative console.

Alias

Specifies the alias for the personal certificate in the keystore.

You enter the alias name for the personal certificate in the keystore when you are creating a certificate. *The alias name is read-only when you view an existing certificate.*

Data type: Text

Version

Specifies the version of the personal certificate. Valid versions include X509 V3, X509 V2, or X509 V1. It is recommended to use X509 V3 certificates.

This field is read-only when you create or view a certificate.

Data type: Text
Default: X509 V3
Range:

Key size

Specifies the key size of the private key that is used by the personal certificate.

When you are creating a certificate you can select the key size from the drop-down list. *This field is read-only when you view a certificate.*

Data type: Integer
Default: 1024
Other valid key sizes: 512, 2048, 4096

Common name

Specifies the common name portion of the distinguished name (DN). It is recommended that this name be the host name of the machine on which the certificate resides. In some cases, the common name is used to login during Secure Socket Layer (SSL) certificate authentication; therefore, in some cases, this name might be used as a user ID for a local operating system registry.

When you create a new certificate you can enter the common name in this field. *This field does not display when you view an existing certificate.*

Data type: Text

Serial number

Identifies the certificate serial number that is generated by the issuer of the certificate. When creating a certificate this field does not appear.

This field is read-only when you view an existing certificate.

Validity period

Specifies the length in days during which the certificate is valid. The default is 365 days. You can enter any number of days you wish.

This field is read-only when you view an existing certificate. This field displays a validity period as a range of days between two dates. For example, Valid from March 16, 2008 to March 16, 2009.

Data type: Text

Organization

You enter the organization portion of the distinguished name. This field is optional.

This field displays only when you create a new certificate.

Data type: Text

Organization unit

Specifies the organization unit portion of the distinguished name. This field is optional.

This field displays only when you create a new certificate.

Data type: Text

Locality

Specifies the locality portion of the distinguished name. This field is optional.

This field displays only when you create a new certificate.

Data type: Text

State/Province

Specifies the state portion of the distinguished name. This field is optional.

This field displays only when you create a new certificate.

Data type: Text

Zip code

Specifies the zip code portion of the distinguished name. This field is optional.

This field displays only when you create a new certificate.

Data type: Integer

Country or region

Select the country portion of the distinguished name from the drop-down list. This field is optional.

This field displays only when you create a new certificate.

Data type: Text
Default: (none)

Refer to <http://www.iso.org/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html> for a list of ISO 3166 country codes.

Validity period

Identifies the length, in days, when the certificate is valid. The default is 365 days.

This field is read-only when you view an existing certificate and shows the start and end dates.

Issued to

Identifies the distinguished name of the entity to which the certificate was issued.

This field is read-only when you view an existing certificate.

Issued by

Identifies the distinguished name of the entity that issued the certificate. When the personal certificate is self-signed, this name is identical to the **Issued to** distinguished name.

This field is read-only when you view an existing certificate.

Fingerprint (SHA Digest)

Identifies the Secure Hash Algorithm (SHA hash) of the certificate, which can be used to verify the certificate's hash at another location, such as the client side of a connection.

This field is read-only when you view an existing certificate.

Signature algorithm

Identifies the algorithm used to sign the certificate.

This field is read-only when you view an existing certificate.

Personal certificate requests collection

Use this page to manage personal certificate requests. Personal certificate requests are temporary place holders for certificates that will be signed by a certificate authority (CA).

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificate requests**.

A private key is generated during the certificate request generation, but only the certificate is sent to the CA. The CA generates a new certificate, signed by the CA. This can be added in the Personal Certificates panel.

Button	Resulting action
New	Creates a personal certificate request that can be given to a certificate authority to complete.

Button	Resulting action
Delete	Deletes a personal certificate request.
Extract	Extracts a personal certificate request. Only one certificate request can be selected at a time.
Query	Queries a personal certificate request. Only one certificate request can be selected at a time.

Note: Any changes made to this panel are permanent.

Key label

Specifies the alias that represents the personal certificate request in the key store.

Requested by

Specifies the Subject distinguished name (DN) that represents the identity of the certificate request.

Personal certificate requests settings

Use this page to create a new certificate request that can be extracted and sent to a certificate authority (CA).

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificates requests**. Then click the **New** button.

Personal certificate requests are temporary place holders for certificates that will be signed by a certificate authority (CA). The private key is generated during the certificate request generation, but only the certificate is sent to the CA. The CA generates a new certificate, signed by the CA.

Note: Any changes made to this panel are permanent.

File for certificate request

Specifies the fully qualified file name from which the certificate request is exported. This portion of the certificate request can be given to the certificate authority to generate the real certificate. After the real certificate is generated, you can perform a "Receive a certificate from a certificate authority" from the personal certificate collection view.

Data type: String

Key label

Specifies the alias that represents the personal certificate request in the key store.

Data type: String

Key size

Specifies the size of the keys that are generated.

Data type: Integer
Default: 1024

Common name

Specifies the name of the entity that the certificate represents. This common name can represent a person, company, or machine. For Web sites, the common name is frequently the DNS host name where the server resides.

Data type: String

Organization

Specifies the organization portion of the distinguished name.

Data type: String

Organizational unit

Specifies the organization unit portion of the distinguished name. This field is optional.

Data type: String

Locality

Specifies the locality portion of the distinguished name. This field is optional.

Data type: String

State or province

Specifies the state portion of the distinguished name. This field is optional.

Data type: String

Zip code

Specifies the zip code portion of the distinguished name. This field is optional.

Data type: Integer

Country or region

Specifies the country portion of the distinguished name.

Data type: String

Refer to <http://www.iso.org/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html> for a list of ISO 3166 country codes.

Extract certificate request

Use this page to extract a certificate request to a file so it can be sent to a certificate authority (CA).

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificate requests > Extract**.

Key label

Specifies the alias that represents the personal certificate request in the key store.

File for certificate request

Specifies the filename where the extracted certificate request is placed.

Data type: Text

Receiving a certificate issued by a certificate authority

When a certificate authority (CA) receives a certificate request, it issues a new certificate that functions as a temporary placeholder for a CA-issued certificate. A keystore receives the certificate from the CA and generates a CA-signed personal certificate that WebSphere Application Server can use for Secure Sockets Layer (SSL) security.

Before you begin

The keystore must contain the certificate request that was created and sent to the CA. Also, the keystore must be able to access the certificate that is returned by the CA.

Note: To receive a certificate by using the wsadmin tool, use the **receiveCertificate** command of the AdminTask object. For more information, see “PersonalCertificateCommands command group for the AdminTask object” on page 1022.

About this task

WebSphere Application Server can receive only those certificates that are generated by a WebSphere Application Server certificate request. It cannot receive certificates that are created with certificate requests from other keystore tools, such as **iKeyman** and **keyTool**.

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates > [keystore]**.
2. Under Additional Properties, click **Personal certificates**.
3. Select a personal certificate.
4. Click **Receive a certificate from a certificate authority**.
5. Type the full path and name of the certificate file.
6. Select a data type from the list.
7. Click **Apply** and **Save**.

Results

The keystore contains a new personal certificate that is issued by a CA. The original certificate request is changed to a personal certificate.

What to do next

The SSL configuration is ready to use the new CA-signed personal certificate.

Export certificate to a keystore file or a managed keystore

Use this page to specify a personal certificate to export to a keystore file or a managed keystore.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > keystore**. Under Additional Properties, click **Personal certificates**. Select a personal certificate using the check box. Then click the **Export** button.

Certificate alias to export:

Displays the name of the certificate that you selected to export on the previous panel.

Data type: Text

Keystore Password:

Type in the password of the keystore to use for the export.

Data type: Text

Alias:

Specifies the alias that the personal certificate is referenced by in the destination keystore.

Data type: Text

Managed key store:

Select this option with the radio button. Then select a keystore from the pull-down list, which is managed by the security configuration, to export the certificate to.

Data type: Drop-down list

Key file name:

Select this option with the radio button. Then type the keystore file name into which the exported certificate is added. If the keystore file name already exists, the exported certificate is added. If the keystore file name does not already exist, a keystore file name is created, and the exported certificate is added.

Data type: Text

Type:

Specifies the type of keystore file. The valid types are listed in the drop-down list.

Data type: Text

Default: PKCS12

Key file password:

Specifies the password that is used to access the keystore file.

Data type: Text

Import certificate from a key file or managed keystore

Use this page to specify a personal certificate to import from a keystore or key file.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} >**

ssl_configuration. Under Related items, click **Key stores and certificates > keystore** . Under Additional Properties, click **Personal certificates**. Select a personal certificate using the check box. Then click the **Import** button.

Managed key store:

Select this option with the radio button. This selection indicates that the keystore that contains the certificate to import is a managed keystore.

Data type: radio button

Get key store aliases:

Clicking this button queries the configuration for the list of keystore aliases for which the certificate will be imported to.

Key store:

Select an alias of the keystore from the pull-down list of managed keystores that are managed by the security configuration. The alias you select identifies the keystore that contains the certificate to import.

Data type: drop-down list

Key store password:

Specifies the password for the keystore to use for import.

Data type: Text

Key store file:

Select this option with the radio button. This selection indicates a keystore file that contains the certificate to import.

Data type: radio button

Key file name:

Specifies the fully qualified path to keystore file that contains the certificate to import.

Data type: Text

Get key file aliases:

Clicking this button, queries the key file for the aliases of all the personal certificates in the keystore from which to choose.

Type:

Specify the type of keystore file. Select a valid type from the drop-down list.

Data type: Text

Key file password:

Type the password that is used to access the keystore file.

Data type: Text

Certificate alias to import:

Specifies the certificate alias identified as the **Key file name** that you want to import into the current keystore.

Data type: Text
Default: (none)

Imported certificate alias:

Specifies the new alias that you want the certificate to be named in the current keystore.

Data type: Text

Receive certificate from CA

Use this page to import your personal certificate from the certificate authority (CA). The imported certificate replaces the temporary certificate associated with the public/private keys in the certificate request that is stored in the key store.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificates > Receive certificate from certificate authority**.

Certificate file name:

Specifies the filename that contains the certificate generated by the certificate authority (CA).

Data type: Text

Data type:

Specifies the format of the file that is either Base64 encoded ASCII data or Binary DER data.

Data type: Text
Default: Base64-encoded ASCII data

Replace a certificate

Use this page to specify two certificates: the first selected certificate is replaced by the second selected certificate. The replace function replaces all the old signer certificates in key stores that are managed throughout the cell with the new signer from the new certificate. The same level of trust that was established with the old certificate is maintained. All places the certificate's alias is referenced in the security configuration will be replaced with the certificate's alias. The alias could be referenced on a security object like the SSL configuration, the dynamic outbound endpoint SSL configuration and key set groups.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties click **Personal certificates**. Select a personal certificate, then click the **Replace** button.

Old certificate

Specifies the certificate that you want to replace.

Data type: Text

Replace with

Specifies the certificate that you want to replace the old certificate.

Data type: Text
Default: (none)

Delete old certificate after replacement

Specifies that you want to delete the old certificate and all associated signer certificates after the new certificate replaces it. If you do not replace the old personal certificate, it might be assigned a new alias name.

Default: Disabled

Delete old signers

Specifies that you want to delete the old signer certificates that are associated with the old certificate after the new signer certificates replace them. If you do not delete the old signer certificates, they might be assigned new alias names.

Default: Disabled

Extracting a signer certificate from a personal certificate

Personal certificates contain a private key and a public key. You can extract the public key, called the *signer certificate*, to a file, then import the certificate into another keystore. The client requires the signer portion of a personal certificate for Security Socket Layer (SSL) communication.

Before you begin

The keystore that contains a personal certificate must already exist.

Note: To extract a signer certificate from a personal certificate using the wsadmin tool, use the **extractCertificate** command of the AdminTask object. For more information, see “PersonalCertificateCommands command group for the AdminTask object” on page 1022.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates > keystore**.
2. Under Additional Properties, click **Personal certificates**.
3. Select a personal certificate.
4. Click **Extract**.
5. Type the full path for the certificate file name. The signer certificate is written to this certificate file.

6. Select a data type from the list.
7. Click **Apply**.

Results

The signer portion of the personal certificate is stored in the file that is provided.

What to do next

This signer can now be imported into other keystores.

Extract certificate

Use this page to extract the signer from the personal certificate and store it in a file. The certificate can be added to a trust store for trust verification. When extracting the signer from a chained personal certificate, the signer at the top level of the chain is extracted.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Personal certificates > Extract**.

Certificate alias to extract

Displays the name of the certificate that you selected for extraction on the previous panel.

Data type: Text

Certificate file name

Specifies the fully qualified path where the certificate file will reside.

Data type: Text

Data type

Specifies the format of the file, which is either Base64-encoded ASCII data or Binary DER data.

Data type: Text
Default: Base64-encoded ASCII data

Extract signer certificate

Use this page to extract a signer certificate from the keystore to a file so that it can be added elsewhere.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Signer certificates**. Select a signer certificate, then click the **Extract** button

File name

Specifies the fully qualified file name where the extracted signer certificate is placed.

Data type: Text

Data type

Specifies the format of the file, which is either Base64 encoded ASCII data or Binary DER data.

Data type:

Text

Retrieving signers using the retrieveSigners utility at the client

The client requires the signer certificates from the server to be able to communicate with WebSphere Application Server. Use the **retrieveSigners** command to get the signer certificate from a server.

Before you begin

The retrieveSigners utility is located in one of the following directories, depending on your operating system:

- **Windows** `profile_root\bin`. For example: `C:\WebSphere\AppServer\profiles\AppSrv01\bin`
- **UNIX** `../profile_root/bin`

In this release, a Java client that does not have access to a stdin console prompt should use the retrieveSigners utility to download the signers from the remote server key store when signers are needed for a Secure Sockets Layer (SSL) handshake. For example, you might interpret the client as failing to respond if an applet client or Java Web Start Client application cannot access the stdin signer exchange prompt. Thus, you must add the WebSphere Java method call **com.ibm.wsspi.ssl.RetrieveSignersHelper.callRetrieveSigners** to your client application to retrieve the signers and to avoid running the retrieveSigners utility manually.

Use the retrieveSigners utility for situations where you cannot verify whether or not the `com.ibm.ssl.enableSignerExchangePrompt=` property is enabled or disabled when the application makes a request. Set the `com.ibm.ssl.enableSignerExchangePrompt=` property to `false` in the `ssl.client.props` file if you cannot see the console.

Alternatively, you can manually create the server key in the client truststore.

About this task

Complete the following steps, as required:

1. Use the **retrieveSigners** command to get the signer certificate from a server. You can find details about the **retrieveSigners** parameters in “Secure installation for client signer retrieval” on page 617.
2. If the client and server are on the same machine, you will need only the `remoteKeyStoreName` and `localKeyStoreName` parameters. The most typical key store to reference on a remote system is `CellDefaultTrustStore` on a network deployed environment and `NodeDefaultTrustStore` on an application server.
3. When retrieving signers from a remote server, add these required connection-related parameters: **-host** *host*, **-port** *port*, **-conntype** {*RMI* | *SOAP*}.
4. Use the **-autoAcceptBootstrapSigner** parameter if you want to enable automation of the signer retrieval. This parameter automatically adds to the server all the signers that are needed to make the connection.

Results

After running, the command displays the SHI-1 digest of the signers added. The output looks similar to the following output:

```
C:\WebSphere\AppServer\profiles\AppSrv01\bin\retrieveSigners.bat
CellDefaultTrustStore ClientDefaultTrustStore
```

```
CWPKI0308I: Adding signer alias "default_signer" to local keystore
             "ClientDefaultTrustStore" with the following SHA digest:
```

Example

The following examples illustrate how to call the `retrieveSigners.bat` file.

To retrieve signers on the same system, enter:

```
profile_root\bin\retrieveSigners.bat CellDefaultTrustStore ClientDefaultTrustStore
```

To retrieve signers on a remote system with a SOAP connection, enter:

```
profile_root\bin\retrieveSigners.bat CellDefaultTrustStore ClientDefaultTrustStore  
-host myRemoteHost -port 8879 -conntype SOAP -autoAcceptBootstrapSigner
```

To retrieve signers on a remote system with an RMI connection, enter:

```
profile_root\bin\retrieveSigners.bat CellDefaultTrustStore ClientDefaultTrustStore  
-host myRemoteHost -port 2809 -conntype RMI -autoAcceptBootstrapSigner
```

To retrieve signers on a remote system that has security enabled, enter:

```
profile_root\bin\retrieveSigners.bat CellDefaultTrustStore ClientDefaultTrustStore  
-host myRemoteHost -port 8879 -conntype SOAP -user testuser -password testuserpwd  
-autoAcceptBootstrapSigner
```

Changing the signer auto-exchange prompt at the client

For clients to communicate with WebSphere Application Server, clients must obtain a signer certificate from the server. Clients can use the **retrieveSigners** command to connect to a server to obtain the appropriate signer. A prompt displays that asks whether or not you want to add a signer to the truststore. If the Secure Sockets Layer (SSL) configuration uses an automated script that might hang, use the prompt to obtain the certificate.

Before you begin

The `com.ibm.ssl.enableSignerExchangePrompt` property in the `profile_home/properties/ssl.client.props` file controls the signer certificate prompt. By default, this property is set to `true`, meaning the prompt is enabled.

About this task

Complete the following steps to disable or enable the signer-exchange prompt at the client:

1. Open the `profile_home/properties/ssl.client.props` file using an editor.
2. Locate the section containing the SSL configuration information for the client that you are working with.
3. Change the value of the `com.ibm.ssl.enableSignerExchangePrompt` property to `false` if you do not want the signer-exchange prompt, or set it to `true` if you want to be prompted.
4. Save and close the file.

Results

When the `com.ibm.ssl.enableSignerExchangePrompt` property is set to `false`, no prompt displays if a signer is not trusted. In this case the SSL handshake fails. Once the proper signer for the connection being made is manually installed in the trust store, the SSL handshake can succeed.

When the `com.ibm.ssl.enableSignerExchangePrompt` property is set to `true`, a signer-exchange prompt displays, and you are asked to accept or reject the certificate. If you accept the certificate, it is installed in the trust store automatically and the handshake succeeds. If you reject the certificate, it does not get installed in the trust store and the handshake fails since the certificate is not trusted.

The prompt looks like the following example:

Example

```
C:\WebSphere\AppServer\profiles\dmgr\bin>serverStatus -all
ADMU0116I: Tool information is being logged in file
           C:\WebSphere\AppServer\profiles\Dmgr\logs\serverStatus.log
ADMU0128I: Starting tool with the dmgr profile
ADMU0503I: Retrieving server status for all servers
ADMU0505I: Servers found in configuration:
ADMU0506I: Server name: dmgr

*** SSL SIGNER EXCHANGE PROMPT ***
SSL signer from target host 192.174.1.5 is not found in truststore
C:/WebSphere/AppServer/profiles/Dmgr/etc/trust.p12.
```

Verify that the digest value matches what is displayed at the server in the following signer information:

```
Subject DN:  CN=hostname.austin.ibm.com, O=IBM, C=US
Issuer DN:   CN=hostname.austin.ibm.com, O=IBM, C=US
Serial number: 1128544457
Expires:     Thu Oct 20 15:34:17 CDT 2006
SHA-1 Digest: 91:A1:A9:2D:F2:7D:70:0F:04:06:73:A3:B4:A4:9C:56:9D:A8:A3:BA
MD5 Digest:  88:72:C5:88:00:1C:A7:FA:D6:EB:04:88:AC:A1:C9:13
```

```
Add signer to the truststore now? (y/n) y
A retry of the request might need to occur.
ADMU0508I: The Application Server "server1" is STARTED.
```

What to do next

Clients can instigate communications for various processes using signer certificates obtained from WebSphere Application Server.

Retrieving signers from a remote SSL port

To perform Secure Sockets Layer (SSL) communication with a server, WebSphere Application Server must retrieve a signer certificate from a secure remote SSL port during the handshake. After the signer certificate is retrieved, you can add the signer certificate to a keystore.

Before you begin

The keystore that is to contain the signer certificate must already exist.

Note: To retrieve a signer certificate from a port using the wsadmin tool, use the **retrieveSignerFromPort** command of the AdminTask object. For more information, see “SignerCertificateCommands command group for the AdminTask object” on page 1043.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > Key stores and certificates > keystore > Signer certificates > Retrieve from port**.
2. Click **Retrieve from port**.
3. Type the host name of the machine on which the signer resides.
4. Type the port location on the host machine on which the signer resides. The port location is not limited to ports on WebSphere Application Server. The ports can include Lightweight Directory Access Protocol (LDAP) ports or ports on any server on which an SSL port is already configured, such as SIB_ENDPOINT_SECURE_ADDRESS.
5. Select an SSL configuration for the outbound connection from the list.

6. Type an alias name for the certificate.
7. Click **Retrieve signer information**. A message window displays information about the retrieved signer certificate, such as: the serial number, issued-to and issued-by identities, SHA hash, and expiration date.
8. Click **Apply**. This action indicates that you accept the credentials of the signer.

Results

The signer certificate that is retrieved from the remote port is stored in the keystore.

What to do next

An SSL configuration or client process that requires an SSL connection to the server can use the retrieved and approved signer certificate.

Retrieve from port

Use this page to retrieve a signer certificate from a remote SSL port. The system connects to the specified remote SSL host and port and receives the signer during the handshake using a trust manager. The signer SHA hash displays for validation and, if approved by an administrator, is added to the currently selected trust store.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Signer certificates**. Then click the **Retrieve from port** button.

Host

Specifies the host name to which you connect when attempting to retrieve the signer certificate from the Secure Sockets Layer (SSL) port.

Data type: Text

Port

Specifies the SSL port to which you connect when attempting to retrieve the signer certificate.

Note: In a network deployment environment, you need to specify the correct secure sockets layer (SSL) port number when attempting to retrieve a signer certificate from a remote SSL port.

- Use the port number associated with the port name, **WC_adminhost_secure**, when retrieving a signer certificate from the deployment manager.
- Use the port number associated with the port name, **CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS**, when retrieving a signer certificate from a node.

All certificates must be in place prior to retrieving them from the deployment manager or from base servers.

Data type: Text

SSL configuration for outbound connection

Specifies the SSL configuration that is used to connect to the previously specified SSL port. This configuration is also the SSL configuration that contains the signer after retrieval. This SSL configuration does not need to have the trusted certificate for the SSL port as it is retrieved during validation and presented here.

Data type: Text
Default: DefaultSSLConfig

Alias

Specifies the certificate alias name that you want to reference the signer in the key store, which is specified in the SSL configuration.

Data type: Text

Retrieved signer information

Specifies the signer certificate information if it is retrieved from the remote host and port.

Serial number

Specifies the certificate serial number that is generated by the issuer of the certificate.

Issued to

Specifies the distinguished name of the entity to which the certificate was issued.

Issued by

Specifies the distinguished name of the entity that issued the certificate. This name is the same as the issued-to distinguished name when the signer certificate is self-signed.

Fingerprint (SHA Digest)

Specifies the Secure Hash Algorithm (SHA hash) of the certificate, which can be used to verify the certificate's hash at another location, such as the client side of a connection.

Expiration

Specifies the expiration date of the retrieved signer certificate for validation purposes.

Adding a signer certificate to a keystore

Signer certificates establish the trust relationship in SSL communication. You can extract the signer part of a personal certificate from a keystore, and then you can add the signer certificate to other keystores.

Before you begin

The keystore that you want to add the signer certificate to must already exist.

Note: To add a signer certificate to a keystore by using the wsadmin tool, use the **addSignerCertificate** command of the AdminTask object. For more information, see `rxml_atsignercert.dita`.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > Inbound | Outbound > SSL_configuration_name > Key stores and certificates**.
2. Select a keystore from the list of keystores.
3. Click **Add signers**.
4. Enter an alias for the signer certificate in the **Alias** field
5. Enter the full path to the signer certificate file in the **File name** field.
6. Select a data type from the list in the **Data type** field.
7. Click **Apply**.

Results

When these steps are completed, the signer from the certificate file is stored in the keystore. You can see the signer in the keystore files list of signer certificates. Use the keystore to establish trust relationships for the SSL configurations.

Add signer certificate settings

Use this page to add a signer certificate in a certificate file to the keystore in the security configuration.

To view this administrative console page, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates > keystore > Signer certificates > Add..**

Alias

Specifies the alias that is used to identify the signer certificate in the keystore.

Data type: String

File name

Specifies the path to the filename where the signer certificate is located.

Data type: String

Data type

Specifies the format of the file, which is either Base64 encoded ASCII data or Binary DER data.

Data type: String

Signer certificates collection

Use this page to manage signer certificates in key stores. Signer certificates are used by Java Secure Socket Extensions (JSSE) to validate certificates sent by the remote side of the connection during a Secure Sockets Layer (SSL) handshake. If a signer does not exist in the trust store that can validate the certificate sent, the handshake fails and generates an "unknown certificate" error.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Signer certificates**.

Button	Resulting action
Add	Adds a new trusted (signer) certificate.
Delete	Deletes an existing signer certificate.
Extract	Extracts a signer certificate from a personal certificate to a file.
Retrieve from port	Makes a test connection to an SSL port and retrieves the signer from the server during the handshake. The information from the certificate will be displayed so you can decide whether to trust it based upon the MD5 and/or SHA hash.

Alias

Specifies the alias for this signer certificate in the key store.

Issued to

Specifies the distinguished name of the entity that requested the certificate.

Fingerprint (SHA digest)

Specifies the Secure Hash Algorithm (SHA hash) of the certificate. This can be used to verify the hash for the certificate at another location, such as the client side of a connection.

Expiration

Specifies the expiration date of the signer certificate for validation purposes.

Signer certificate settings

Use this page to verify the general properties of the selected signer certificate.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates > key store**. Under Additional Properties, click **Signer certificates**. Then click on a signer certificate.

Alias

Specifies the alias for this signer certificate in the key store.

Version

Specifies the version of the personal certificate. Valid versions include X509 V3, X509 V2, or X509 V1.

Key size

Specifies the key size of the public key used by the signer certificate.

Serial number

Specifies the certificate serial number that is generated by the issuer of the certificate.

Validity period

Specifies the begin and end dates of the certificate.

Issued to

Specifies the distinguished name of the entity that requested the certificate.

Issued by

Specifies the distinguished name of the entity that issued the certificate. This name is the same as the issued-to distinguished name when the signer certificate is self-signed.

Fingerprint (SHA Digest)

Specifies the Secure Hash Algorithm (SHA) hash of the certificate, which can be used to verify the hash for the certificate at another location such as the client side of a connection.

Signature algorithm

Specifies the algorithm that is used to sign the certificate.

Adding a signer certificate to the default signers keystore

Signer certificates are added to a keystore on the client side of an SSL communication to establish trust with the server. There is common practice for keystores to have trust established when they are created. The **DmgrDefaultSignersStore** on a deployment manager and the **NodeDefaultSignersStore** on a stand alone application server are created to hold signer certificates used to establish trust by default in newly create keystores.

Before you begin

The default signers key store is created during profile creation and contains the signer certificate of the server default root certificate. Additional signer certificates can be added to the default signers key store at any time. Anytime a keystore is created using the admin console or by using the **createKeyStore** AdminTask object in scripting, all signer certificates from the default signer store are added to the newly created keystore.

Note:

- To add a signer certificate to a default signer keystore by using the wsadmin tool, use the **addSignerCertificate** command of the AdminTask object.
- To create a new keystore by using the wsadmin tool, use the **createKeyStore** command of the AdminTask object.
- To extract the signer from a personal certificate using the wsadmin tool, use the **extractCertificate** of the AdminTask object.
- To exchange a signer certificate using the wsadmin tool, use the **KeyStoreCommands** command group for the AdminTask object.

For more information, see `rxml_atsignercert.dita` and `rxml_atkeystore.dita`

1. If the certificate is in a certificate file, it can be added to the default signer keystore using the administrative console.
 - a. Click **Security > SSL certificate and key management**.
 - b. Under Related Items, click **Key stores and certificates**.
 - c. Select Default signers keystore under KeyStore Usages. A panel displaying a list of keystores appears.
 - d. Click on **DmgrDefaultSignersStore**.
 - e. Under Additional Properties, click **Signer certificates**.
 - f. Click **Add**.
 - g. Enter an alias in the alias box, a path to the certificate file in the filename box, and an asterisk (*). Select the format of the certificate file from the pull down list in the “Data type” box.
 - h. Click **Apply** then **Save**.

Note: You can also perform this addition using the AdminTask, **addSignerCertificate**.

2. If the signer certificate form of a personal certificate needs to be added to default signers keystore, you can extract the signer from the personal certificate to a certificate file or the signer can be extracted directly to the default signers keystore. To extract a signer certificate from a personal certificate to a certificate file,
 - a. Click **Security > SSL certificate and key management**.
 - b. Under Related Items, click **Key stores and certificates**.
 - c. Select All under Keystore Usages. A panel displaying a list of keystores appears.
 - d. Click on the keystore name
 - e. Under Additional Properties, click **Personal certificates**.
 - f. Select a personal certificate.
 - g. Click **Extract**.
 - h. Enter the path to the certificate file in “Certificate file name” box and select a format type from the pull down list in “Data type” box
 - i. Click **Apply** then **Save**.
 - j. The signer can be added to the default signers keystore by following step 1.

Note: You can also extract the signer from a personal certificate using scripting and the AdminTask `extractCertificate`.

3. To extract a signer certificate to the default signers keystore, an exchange of the signer certificate can be performed from the administrative console.
 - a. Click **Security > SSL certificate and key management**
 - b. Under Related Items, click **Key stores and certificates**.
 - c. Select All under Keystore Usages. A panel displaying a list of keystores appears.
 - d. Click on the default signers keystore and the keystore that contains the personal certificate whose signer certificate is needed.
 - e. Click **Exchange Signers**.
 - f. Select the personal certificate whose signer is needed.
 - g. Click **Add**.
 - h. Click **Apply** then **Save**.

Note: You can also perform the exchange using the AdminTask, `exchangeSigner`.

Results

When these steps are completed, the signer from the certificate file is stored in the default signers keystore. You can see the signer in the keystore files list of signer certificates.

What to do next

The new keystore will contain the default signers that were added to the default signers keystore.

Exchanging signer certificates

To establish trust relationships, you can exchange signer certificates between keystores. When you exchange signer certificates, you are extracting a personal certificate from one keystore and adding it to another keystore as a *signer certificate*.

Before you begin

To exchange signer certificates, there must be two keystores.

About this task

Complete the following steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key stores and certificates**.
2. Select two keystores from the list of keystores.
3. Click **Exchange signers**.
4. Select any of the certificates in the first personal certificates list, and click **Add**. After adding, the signer part of the selected personal certificate appears in the other (second) keystore signers list.
5. Select any of the certificates in the second personal certificates list, and click **Add**. After adding, the signer part of the selected personal certificate appears in the other (first) keystore signers list.
6. Optional: If you need to remove any of the certificates from either of the signers lists, highlight one or more of the certificates, and click **Remove**.
7. Click **Apply** and **Save**.

Results

The signer certificate appears in the list for each keystore.

What to do next

The extracted signer certificate is available to both keystores during the connection handshake.

Keystores and certificates exchange signers

Use this page to extract the signer part of a personal certificate from one keystore and add it to another keystore as a signer certificate. Signer certificates can also be listed, and they will be added to the other keystore as well.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key stores and certificates** then select two key stores to exchange and click the **Exchange signers**.

Note: Any changes made to this panel are permanent.

[keystore] personal certificates

Specifies the personal certificates and signer certificates that are currently stored in the specified keystore.

Press and hold the **Ctrl** key to select more than one item from the list.

Data type: Text

[keystore] signers

Specifies the trusted signer certificates that are currently stored in the specified keystore and selected for the exchange.

Press and hold the **Ctrl** key to select more than one item from the list.

Data type: Text

Add

Specifies to extract the signer from the selected personal certificate in the keystore list on the left and add it to the signers list of the keystore on the right.

After the certificate is added, it no longer displays in the left-hand list. The personal certificate is still in the keystore, but it is no longer selectable

Remove

Specifies to remove a selected signer from the signers list of the keystore on the right. The removed certificate displays in the keystore list on the left.

Configuring certificate expiration monitoring

When certificates expire, they can no longer be used by the system. WebSphere Application Server provides a utility to monitor certificates that are close to expiration or have already expired. You can schedule certificate monitoring, or you can request certificate monitoring on demand. You can also configure options for deleting expired certificates and for recreating certificates.

Before you begin

WebSphere Application Server notifies you when a certificate is about to expire. Complete the information required for notification messaging in “Notifications” on page 714.

About this task

Complete the following configuration steps in the administrative console:

1. Click **Security > SSL certificate and key management > Manage certificate expiration**.
2. Type a number for the number of days threshold in the **Expiration notification threshold** field. WebSphere Application Server issues an expiration warning n number of days before expiration.
3. Select or check one or more of the following options:
 - **Expiration check notification**. Select the method from the list that you want to use to receive your notification.
 - **Automatically replace expiring self-signed certificates**. If you do not want to recreate the self-signed certificate, clear the check box.

Note: When using writable System Authorization Facility (SAF) keyrings in your configuration, the certificate expiration monitor does not replace expired certificates in the writable SAF keyrings, but only provides a notification of the expiration.

 - **Delete expiring certificates and signers after replacement**. If you do not want to delete the expired certificates and signers, clear the check box.
 - **Enable checking**. If you do not want to have certificate monitoring enabled, clear the check box.
4. Enter the time of day when you want certificate monitoring to take place to schedule the running of the certificate expiration monitor.
5. Select one of the following options:
 - **Check by calendar**. For **Weekday**, enter the day of week that you want to run the certificate expiration monitor. For **Repeat Interval**, specify the frequency to run the certificate monitor.
 - **Check by number of days**. Enter a number for how frequently the monitor runs, in number of days.
6. Click **Apply**.

Results

After completing the settings, a certificate expiration monitor object and a schedule are set up in the configuration. The certificate expiration monitor runs according to the configurations options that you configured.

What to do next

You can generate reports that state which certificates have expired. The reports identify the notifications of certificate replacements and deletions. The report is sent according to the notification option that you specified.

Manage certificate expiration settings

Use this page to configure the certificate expiration monitor.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage certificate expiration**.

Note: To see the changes to the Expiration checking fields, you must click **Apply**.

Start now

Specifies to start certificate monitoring. When the monitor runs, it visits all the key stores and checks to see if they are within certificate expiration range. If you set the option to delete or replace expired certificates, you can run these operations immediately by pressing **Start now**.

Expiration notification threshold

Specifies the period of time that occurs chronologically just before the expiration day of the certificate, within which, if the ExpirationMonitor thread runs, and **Automatically replace expiring self-signed and chained certificates** is enabled, a new self-signed or chained certificate is generated. By default, the replacement period for the certificate is 60 days in length or less as defined in the daysBeforeNotification property.

There is a pre-notification period where the certificate is added to the notification list but not touched for 90 days prior to the 60 days. By default, this pre-notification period is 90 days in length as defined in the com.ibm.ws.security.expirationMonitorNotificationPeriod property.

Data type:	Integer
Default:	60 days or less

Enable checking

Specifies the certificate monitor is active and will run as scheduled.

Scheduled time of day to check for expired certificates

Specifies the scheduled time that the system checks for expired certificates.

You can type the scheduled time in hours and minutes, specify either A.M. or P.M., or 24-hour.

Data type	Integer
Default:	0, 0
Range:	1–12, 0–59

Check by calendar

Indicates that you want to schedule a specific day of the week on which the expiration monitor runs. For example, it might run on Sunday.

Default:	Disabled
-----------------	----------

Weekday

Specifies the day of the week on which the expiration monitor runs if **Check on a specific day** is selected.

Default:	Sunday
Range:	Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday

Repeat interval

Specifies the period of time between each schedule time to check for expired certificates or the interval between schedule checks.

Default:	Daily
Range:	Daily, Weekly

Check by number of days

Specifies that you want to schedule a specific number of days between each run of the expiration monitor. The day of the week on which this occurs is not counted. For example, if you set the interval to check for expired certificates every seven days, the expiration monitor runs on day eight.

Default: Disabled

Next start date

Specifies the date for the next scheduled check. This allows the deployment manager to be stopped and restarted without resetting the date.

Expiration check notification

Specifies the notification type (such as e-mail or System Out) when an expiration monitor runs.

Default:

Automatically replace expiring self-signed certificates and chained certificates

Specifies a new self-signed certificate or chained certificate be generated using the same certificate information if the expiration notification threshold is reached. The old certificate is replaced and uses the same alias. All old signers are managed by the key store configuration are also replaced. The system only replaces self-signed certificates.

Note: This checkbox is only applicable when using file based keystores.

Default: Enabled

Delete expiring certificates and signers after replacement

Specifies whether to completely remove old, self-signed certificates from the key store during a replace operation or leave them there under a renamed alias. If an old certificate is not deleted, the system renames the alias so that the new certificate can use the old alias, which might be referenced elsewhere in the configuration.

Note: This checkbox is only applicable when using file based keystores.

Default: Enabled

Notifications

Use this page to specify the generic notification definitions that are used in certificate expiration monitors.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage certificate expiration**. Under Related items, click **Notifications**.

Button	Resulting action
New	Adds a notification. The notification configures how the expiration monitor notifies the administrator of certificates that will expire within the specified threshold.
Delete	Deletes an existing notification.

Notification name

Specifies the notification name.

Message log

Specifies that this configuration intends to log certificate expiration information to the message log file.

Send E-mail

Specifies that this configuration intends to send certificate expiration information to the list of users in the e-mail list.

List of E-mail Addresses

Specifies the e-mail addresses that are sent notifications when certificates fall within the expiration threshold. You must specify the SMTP server for each e-mail address. If an e-mail address is not specified, by default the application server assumes that the SMTP server is "smtp-server." For example, if you type *name@domain*, the SMTP server will be *smtp-server.domain*.

Notifications settings

Use this page to set properties for new notifications used in certificate expiration monitors or for security audit subsystem failures.

To view this administrative console page perform one of the following:

- Click **Security > SSL certificate and key management > Manage certificate expiration > Notifications > New**.
- Click **Security → Security auditing → Audit monitor → New**

Notification name

Specifies the name of the notification configuration.

Data type: Text

Message log

Specifies that this configuration will log the notification to a message log file.

Default: Disabled

Email sent to notification list

Specifies that this configuration send a notification as an e-mail to the e-mail list.

Default: Disabled

Email address to add

Specifies the e-mail addresses that are sent notifications. You must specify the SMTP server for each e-mail address. If an e-mail address is not specified, by default the application server assumes that the SMTP server is "smtp-server." For example, if you type *name@domain*, the SMTP server will be *smtp-server.domain*.

Data type: Text (format as valid Internet mail address)

Add

Adds the e-mail address to the right-hand list.

Remove

Removes the e-mail address from the right-hand list.

Outgoing mail (SMTP) server

Specifies the SMTP server to be used with the e-mail address. If none is specified, the e-mail realm will be used.

Key management for cryptographic uses

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

The key management infrastructure is based on two key configuration object types: key sets and key set groups. WebSphere Application Server uses a key set to manage instances of keys of the same type. You can configure a key set to generate a single key or a key pair, depending on the key or key pair generator class. A key set group manages one or more key sets and enables you to configure and generate different key types at the same time. For example, if your application needs both a secret key and key pair for cryptographic operations, you can configure two key sets, one for the key pair and one for the secret key that the key set group manages. The key set group controls the auto-generation characteristics of the keys, including the schedule. The framework can automatically generate keys on a scheduled basis, such as on a particular day of the week and time of day, so that key generation is done during off-peak hours.

Figure 1 shows an example of a key set group that is configured to manage two key sets: key set 1 and key set 2.

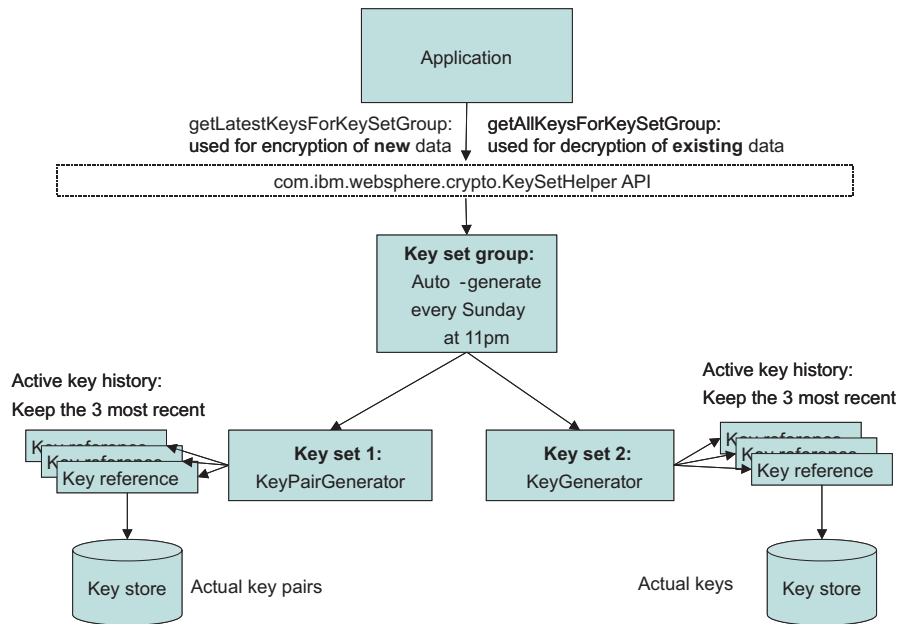


Figure 1: Key Management Infrastructure

Figure 19.

Key set 1 generates key pairs. Key set 2 generates secret keys. The application needs both types of keys for its cryptographic operations, signing and encryption, on data. The keys for each key set need to be generated in tandem. The application stores the key set group name with the encrypted data. The key set group generates a new set of keys every Sunday night at 11 P.M.. The application maintains key generation data for two weeks.

Creating a key set configuration

You can use key sets to manage multiple instances of cryptographic keys. WebSphere Application Server uses keys to encrypt or sign outbound data, and decrypt or verify inbound data during cryptographic operations.

Before you begin

You must have write-access to the keystore that will contain the keys after you generate them from a key set. However, if you want to generate keys outside of WebSphere Application Server, you can reference the keys from a read-only keystore that contains a secret key that you can access when you generate the keys. If you are creating a key pair using an X509Certificate and a PrivateKey object, see “Example: Developing a key or key pair generation class for automated key generation” on page 724.

About this task

Complete the following steps in the administrative console:

1. Decide whether you want to create the key set at the cell scope or below the cell scope at the node, server, or cluster, for example:
 - To create a key set at the cell scope, click **Security > SSL certificate and key management > Key sets**.
 - To create a key set at a scope below the cell level, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration > Key sets**.
2. Click **New** to create a new key set.
3. Type a key set name. For example, CellmyKey.
4. Type a key alias prefix name. For example, myKey. This field specifies the prefix for the key alias when the new key is generated and stored in the keystore. Following the prefix is the key reference version number, for example, 2, so that the full key alias name would be myKey_2. If the key reference already has a specified alias for a key that exists in the keystore, then WebSphere Application Server ignores this field.
5. Type a key password. The key password protects the key in the keystore. This password is ignored by WebSphere Application Server if you already specified a password for the key alias reference. To check for a key reference password, click **Active key history** under Additional Properties. The key reference password protects keys that are generated by a key generator class.
6. Type the password again to confirm it.
7. Optional: Type the key generator class name. For example, com.ibm.ws.security.ltpa.LTPAKeyGenerator. The class name generates keys. If the class implements com.ibm.websphere.crypto.KeyGenerator, then a getKey method returns a java.security.Key object that is set in the keystore using the setKey method without a certificate chain. If the class implements com.ibm.websphere.crypto.KeyPairGenerator, then a getKeyPair method returns a com.ibm.websphere.crypto.KeyPair object that contains either a java.security.PublicKey and java.security.PrivateKey or a java.security.cert.Certificate and a java.security.PrivateKey object. The key generator class and the KeySetHelper API specify the details of the keys that are generated.
8. Optional: Select **Delete key references that are beyond the maximum number of keys** if you do not want old keys saved in the keystore after WebSphere Application Server removes their references from the Active key history listing. The Active key history lists the keys that the KeySetHelper API is currently tracking. The number of keys in the list is equal to the number of keys that you specify in **Maximum number of keys referenced**.
9. Type a numeric value for the maximum number of keys referenced. For example, if you type 3 and select **Delete key references that are beyond the maximum number of keys**, the fourth key version generation automatically triggers WebSphere Application Server to delete the first key version from the keystore. If you choose not to delete the old keys, they do not display in the Active key history list but instead remain in the keystore where you can remove them manually.
10. Select a keystore from the drop-down list.
 - Select a JCEKS keystore if you are storing a secret key.
 - Select any keystore if you are storing a key pair with an X509Certificate and PrivateKey object.
11. Optional: Select **Generates key pair** if your key generator class name implements the com.ibm.websphere.crypto.KeyPairGenerator interface instead of the com.ibm.websphere.crypto.KeyGenerator interface. This option designates that the key references a key pair instead of a single key. A key pair contains both a public key and a private key. The WebSphere Application Server run time determines whether or not key pairs are stored and loaded differently than single keys.
12. Optional: Click **Apply** if you want to select **Active key history** under Additional Properties to add alias references or generate more keys.
 - a. Click **Active key history**.
 - b. Click **Add key alias reference** if you are not using the key generator class name to add key alias references to the keys that already exist in the keystore. Use this option to retrieve the keys from a read-only keystore without the key set generating them.

- c. Type an alias reference.
 - d. Click **Generate key** if you want to generate a key using the class name that you defined in the key sets panel. Each new key increments numerically, for example, myAlias_2.
 - e. Click **Apply**.
13. Click the key set name in the navigation path at the top of the panel.
 14. Click **OK** and **Save**.

Results

You have created a key set that you can manage using the **Active key history** link. You can generate keys manually to associate them with specified key sets.

What to do next

After you generate new keys from a key set, you can access them programmatically using the `com.ibm.websphere.crypto.KeySetHelper` API. You must have Java 2 Security permissions, if enabled, to access keys in key sets. Specify the key set name within the fine-grained permissions, as in the following code sample: `WebSphereRuntimePermission "getKeySets.keySetName".` For more information, see “Example: Retrieving the generated keys from a key set group” on page 723. To generate multiple key types at the same time or to schedule the key generation on a specific schedule, see “Creating a key set group configuration” on page 722.

Active key history collection

Use this page to manage key alias references.

To view this administrative console page, click **Security** → **SSL certificate and key management** → **Manage endpoint security configurations** > **{Inbound | Outbound}** → **ssl_configuration** → **Key Sets** → **key set** → **Active key history**.

Button	Resulting action
Alias reference	Adds a reference to a key that already exists in a key store. If a key generation class is configured, the references are added automatically during generation and do not need to be added manually.
Delete	Deletes an existing key reference. This action does not delete the key in the keystore.
Generate key	Generates a key. The button is displayed only if a generator class name is specified for the key set, and the selected key store is editable.

Key alias reference

Specifies the name of the alias as it appears in the keystore.

Add key alias reference settings

Use this page to access key alias reference information.

To view this administrative console page, Under Configuration settings, click **Manage endpoint security configurations** > **{Inbound | Outbound}** > **ssl_configuration**. Under Related items, click **Key Sets** > **key set**. Under Additional Properties, click **Active key history** then click the **Add key alias reference** button.

Alias reference

Specifies the name of the alias as it appears in the key store.

Data type: Text

Password

Specifies the key password to get access to the key. This password is enforced by the keystore for that specific key. If the key does not have a password, this field can be left blank.

Data type: Text

Confirm password

Confirms the password entered in the previous field.

Data type: Text

Key sets collection

Use this page to manage key sets, which control a set of key instances of the same type for use in cryptographic operations. The keys can either be generated using a custom class or reference keys that already exist in a keystore.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key sets**.

Button	Resulting action
New	Adds a new key set.
Delete	Deletes an existing key set. Make sure the key set is not referenced by a key set group before deleting it.

Key set name

Specifies the key set name that is used to select the key set from a key set group and from runtime application programming interfaces (API).

Key store

Specifies the key store that contains the keys for storage, retrieval, or both.

Key alias prefix name

Specifies the prefix for the key alias when a new key is generated and stored in a key store. The rest of the key alias comes from the key reference version number.

For example, if the alias prefix is `mykey` and the key reference version is 2, the keystore references the key using alias `mykey_2`. If the key reference already has a specified alias for a key already existing in the keystore, this field is ignored.

Key sets settings

Use this page to set the properties for a new key set.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key sets**. Then click the **New** button.

Key set name

Specifies the key set name that is used to select the key set from a key set group and from runtime application programming interfaces (API).

Data type: Text

Key alias prefix name

Specifies the prefix for the key alias when a new key is generated and stored in a keystore. The rest of the key alias comes from the key reference version number. For example, if the alias prefix is mykey and the key reference version is 2, the keystore references the key using mykey_2. If the key reference already has a specified alias for a key already existing in the keystore, this field is ignored.

Data type: Text

Key password

Specifies the password used to protect the key in the keystore. If a password is specified in the key reference as well, this password is ignored. This password is used for keys that get generated by a key generator class.

Data type: Text

Confirm password

Specifies the same password again to confirm it was entered correctly the first time.

Data type: Text

Key generator class name

Specifies the class name that generates keys. If the class implements `com.ibm.websphere.crypto.KeyGenerator`, then a `getKey()` method should return a `java.security.Key` object that is set in the key store using the `setKey` method without a certificate chain. The key store type associated with the key set must support storing keys without certificates, such as JCEKS.

Data type: Text

If the class implements `com.ibm.websphere.crypto.KeyPairGenerator`, then a `getKeyPair()` method should return a `com.ibm.websphere.crypto.KeyPair` object containing either a `java.security.PublicKey` and `java.security.PrivateKey`, or a `java.security.cert.Certificate[]` and a `java.security.PrivateKey`. The key generator class and the caller of the `KeySetHelper` API should know the details of the keys that are generated. This framework does not need to understand the key algorithms and lengths.

Delete key references that are beyond the maximum number of keys:

Specifies that the keys are deleted from the keystore at the same time that the key reference is deleted. The server deletes the older key references as the Maximum number of keys referenced value is exceeded.

Maximum number of keys referenced

Specifies the maximum number of key instances that are returned when keys from this key set are requested. The oldest key reference gets removed whenever a new key reference gets generated after the maximum has been reached.

Data type: Integer

Default: 3

Key store

Specifies the key store that contains the keys for storage, retrieval, or both.

Data type:

Text

Generates key pair

Specifies that a key references a key pair instead of a key. The key pair contains both a public key and a private key.

Creating a key set group configuration

A key set group manages one or more key sets. WebSphere Application Server uses key set groups to automatically generate cryptographic keys or multiple synchronized key sets.

About this task

Complete the following steps in the administrative console:

1. Decide whether you want to create the key set group at the cell scope or below the cell scope at the node, server, or cluster, for example.
 - To create a key set group at the cell scope, click **Security > SSL certificate and key management > Key set groups**.
 - To create a key set group at a scope below the cell level, click **Security > SSL certificate and key management > Manage endpoint security configurations > {Inbound | Outbound} > SSL_configuration > Key set groups**.
2. You can choose to generate a key for an existing key set group, delete an existing key set group, or create a new key set group.
 - To generate a key for an existing key set group, select a key set group from the list of existing key set groups, and click **Generate keys**. You have generated a new key for each key set in the selected group.
 - To delete an existing key set group, select a key set group from the list of existing key set groups, and click **Delete**. You have deleted the key set group.
 - To create a new key set group, go to step 3.

Note: Do not delete the cell or node LTPAKeySetGroup, which is used by the Lightweight Third Party Authentication (LPTA) mechanism.

3. Click **New** to create a new key set group.
4. Type a key set group name. You can reference this name by using the `com.ibm.websphere.crypto.KeySetHelper` API to retrieve the managed keys from an application.
5. Select one or more key sets from the Key sets list.

Note: If the key set(s) you want is not listed, make sure that it was created at the same scope or a higher scope than where you are creating the new key set group.

6. Click **Add** to add the selected key set(s) to the new key set group.
7. Select **Automatically generate keys** to generate the new keys on a schedule. If you decide to generate keys automatically, then you must specify a scheduled time of day.
8. Specify the scheduled time to generate keys automatically in hours and minutes, A.M. or P.M., or every 24 hours.
9. You can choose to generate new keys on a specific day or at an interval.
 - Select **Generate on a specific day**. Select a day of the week from the drop-down list, and type a repeat interval number for the number of days between each key generation. This choice enables you to schedule key generation when your systems are least busy.
 - Select **Generate at an interval**. Type a repeat interval number for the number of days between each key generation. This choice enables you to schedule key generation more frequently than once a week.

Note: The **Next start date** is a read-only field that specifies the date for the next scheduled generation. You can stop and restart the deployment manager or base application server without resetting this date. If you do not see the next start date appear after changing the configuration, click **OK** to save it, then check that the next start date displays.

10. Click **Save**.

Results

You have created a new key set group to manage key sets and key generation on a schedule.

What to do next

After you generate new keys from a key set, you can access them programmatically using the `com.ibm.websphere.crypto.KeySetHelper` API. You must have Java 2 Security permissions, if enabled, to access keys in key sets. Specify the key set name within the fine-grained permissions, as in the following code sample: `WebSphereRuntimePermission "getKeySets.keySetName".` For more information, see “Example: Retrieving the generated keys from a key set group.”

Example: Retrieving the generated keys from a key set group

This example shows how applications can use the `com.ibm.websphere.crypto.KeySetHelper` API to retrieve managed keys from the `KeySet` or `KeySetGroup` configurations. Use the `com.ibm.websphere.crypto.KeySetHelper` API to get either the latest set of keys or all the keys in the `KeySet` or `KeySetGroup` object.

Use the latest keys when performing any new cryptographic operations. All of the other keys that are defined in the `KeySet` or `KeySetGroup` object are for the validation of previously performed cryptographic operations.

The following example uses a method that an application might use to initialize the keys in the associated `KeySetGroup` object. The application might want to store the keys in two separate maps, one for generation and one for validation. Refer to the API documentation for `KeySetHelper` API to determine which Java 2 Security requirements are required.

```
/**
 * Initializes the primary and secondary Maps used for initializing the keys.
 */

public void initializeKeySetGroupKeys() throws com.ibm.websphere.crypto.KeyException
{
    java.util.Map generationKeys = null;
    java.util.Map validationKeys = null;

    PublicKey tempPublicKey = null;
    PrivateKey tempPrivateKey = null;
    byte[] tempSharedKey = null;

    keySetGroupName = "ApplicationKeySetGroup";
    com.ibm.websphere.crypto.KeySetHelper ksh = com.ibm.websphere.crypto.KeySetHelper.getInstance();
    generationKeys = ksh.getLatestKeysForKeySetGroup(keySetGroupName);

    /**
     * Latest keys: {
     * KeyPair_3=com.ibm.websphere.crypto.KeyPair@64ec64ec,
     * Secret_3=javax.crypto.spec.SecretKeySpec@fffe8aa7
     * }
     */
    /**/

    if (generationKeys != null)
    {
        Iterator iKeySet = generationKeys.keySet().iterator();
```

```

while (iKeySet.hasNext())
{
    String keyAlias = (String)iKeySet.next();

    Object key = generationKeys.get(keyAlias);

    if (key instanceof java.security.Key)
    {
        tempSharedKey = ((java.security.Key)key).getEncoded();
    }
    else if (key instanceof com.ibm.websphere.crypto.KeyPair)
    {
        java.security.Key publicKeyAsSecret =
((com.ibm.websphere.crypto.KeyPair)key).getPublicKey();
        tempPublicKey = new PublicKey(publicKeyAsSecret.getEncoded());
        java.security.Key privateKeyAsSecret =
((com.ibm.websphere.crypto.KeyPair)key).getPrivateKey();
        tempPrivateKey = new PrivateKey(privateKeyAsSecret.getEncoded());
    }
}

// save these for use later, if necessary
validationKeys = ksh.getAllKeysForKeySetGroup(keySetGroupName);

/**
 * All keys: {
 * version_1=
 * {Secret_1=javax.crypto.spec.SecretKeySpec@178cf,
 * KeyPair_1=com.ibm.websphere.crypto.KeyPair@1c121c12},
 * version_2=
 * {Secret_2=javax.crypto.spec.SecretKeySpec@17a77,
 * KeyPair_2=com.ibm.websphere.crypto.KeyPair@182e182e},
 * version_3=
 * {Secret_3=javax.crypto.spec.SecretKeySpec@fffe8aa7,
 * KeyPair_3=com.ibm.websphere.crypto.KeyPair@4da04da0}
 * }
 */
}
else
{
    throw new com.ibm.websphere.crypto.KeyException("Could not generateKeys.");
}
}

```

Example: Developing a key or key pair generation class for automated key generation

A class that generates keys for cryptographic operations can be created automatically. With this capability, the key management infrastructure can maintain a list of keys for a predefined key set, and applications can access these keys.

You can schedule new key generation at predefined frequencies. Remember that key generation frequency affects the security of your data. For example, for persistent data, you might schedule key generation less frequently than for real time communications, which require that the keys be generated more often as old keys expire.

When you develop a key generation class, decide if you are creating a shared key or a key pair because this decision determines the interface you must use.

If you are developing shared keys, refer to the following example, which uses the KeyGenerator class to implement the com.ibm.websphere.crypto.KeyGenerator interface. The interface returns a java.security.Key key, which is stored as a SecretKey in a JCEKS keystore type. You can use any other keystore type that supports storing secret keys.

```

package com.ibm.test;

import java.util.*;
import com.ibm.ws.ssl.core.*;
import com.ibm.ws.ssl.config.*;
import com.ibm.websphere.crypto.KeyException;

public class KeyGenerator implements com.ibm.websphere.crypto.KeyGenerator
{
    private java.util.Properties customProperties = null;
    private java.security.Key secretKey = null;

    public KeyGenerator()
    {
    }

    /**
     * This method is called to pass any custom properties configured with
     * the KeySet to the implementation of this interface.
     *
     * @param java.util.Properties
     */
    public void init (java.util.Properties customProps)
    {
        customProperties = customProps;
    }

    /**
     * This method is called whenever a key needs to be generated either
     * from the schedule or manually requested. The key is stored in the
     * KeyStore referenced by the configured KeySet that contains the
     * keyGenerationClass implementing this interface. The implementation of
     * this interface manages the key type. The user of the KeySet
     * must know the type that is returned by this keyGenerationClass.
     *
     * @return java.security.Key
     * @throws com.ibm.websphere.crypto.KeyException
     */
    public java.security.Key generateKey () throws KeyException
    {
        try
        {
            // Assume generate3DESKey is there to create the key.
            byte[] tripleDESKey = generate3DESKey();
            secretKey = new javax.crypto.spec.SecretKeySpec(tripleDESKey, 0, 24, "3DES");

            if (secretKey != null)
            {
                return secretKey;
            }
            else
            {
                throw new com.ibm.websphere.crypto.KeyException ("Key could not be generated.");
            }
        }
        catch (Exception e)
        {
            e.printStackTrace(); // handle exception
        }
    }
}

```

If you are developing a key pair, refer to the following example, which uses the KeyPairGenerator class to implement the com.ibm.websphere.crypto.KeyPairGenerator interface.

```

package com.ibm.test;

import java.util.*;
import javax.crypto.spec.SecretKeySpec;
import com.ibm.websphere.crypto.KeyException;

```



```

/**
 * This implementation defines the method to generate a java.security.KeyPair.
 * When a keyGeneration class implements this method, the generateKeyPair method
 * is called and a KeyPair is stored in the keystore. The isKeyPair
 * attribute is ignored since the KeyGenerationClass is an
 * implementation of KeyPairGenerator. The isKeyPair attributes is for when
 * the keys already exist in a KeyStore, and are just read (not
 * generating them).
 *
 * @author IBM Corporation
 * @version WebSphere Application Server 6.1
 * @since WebSphere Application Server 6.1
 */
public class KeyPairGenerator implements com.ibm.websphere.crypto.KeyPairGenerator
{
    private java.util.Properties customProperties = null;

    public KeyPairGenerator()
    {
    }

    /**
     * This method is called to pass any custom properties configured with
     * the KeySet to the implementation of this interface.
     *
     * @param java.util.Properties
     */
    public void init (java.util.Properties customProps)
    {
        customProperties = customProps;
    }

    /**
     * This method is called whenever a key needs to be generated either
     * from the schedule or manually requested and isKeyPair=true in the KeySet
     * configuration. The key is stored in the KeyStore referenced by
     * the configured KeySet which contains the keyGenerationClass implementing
     * this interface. The implementation of this interface manages the
     * type of the key. The user of the KeySet must know the type that
     * is returned by this keyGenerationClass.
     *
     * @return com.ibm.websphere.crypto.KeyPair
     * @throws com.ibm.websphere.crypto.KeyException
     */
    public com.ibm.websphere.crypto.KeyPair generateKeyPair () throws KeyException
    {
        try
        {
            java.security.KeyPair keyPair = generateKeyPair();

            // Store as SecretKeySpec
            if (keyPair != null)
            {
                java.security.PrivateKey privKey = keyPair.getPrivate();
                java.security.PublicKey pubKey = keyPair.getPublic();

                SecretKeySpec publicKeyAsSecretKey = new SecretKeySpec
                    (pubKey.getEncoded(), "RSA_PUBLIC");
                SecretKeySpec privateKeyAsSecretKey = new SecretKeySpec
                    (privKey.getEncoded(), "RSA_PRIVATE");

                com.ibm.websphere.crypto.KeyPair pair = new
                com.ibm.websphere.crypto.KeyPair(publicKeyAsSecretKey, privateKeyAsSecretKey);
                return pair;
            }
            else
            {
                throw new com.ibm.websphere.crypto.KeyException ("Key pair could
                not be generated.");
            }
        }
        catch (Exception e)
    }
}

```

```

    {
        e.printStackTrace(); // handle exception
    }
}

```

This interface returns a `com.ibm.websphere.crypto.KeyPair` key pair, which can contain either a `X509Certificate` and `PrivateKey` object or `PublicKey` and `PrivateKey` objects. If the `com.ibm.websphere.crypto.KeyPair` interface contains a `X509Certificate` and `PrivateKey` object, the certificate and private key are stored in the keystore. Consequently, they can use any `KeyStore` type.

If the `com.ibm.websphere.crypto.KeyPair` interface contains `PublicKey` and `PrivateKey` objects, you must convert the encoded values to the `SecretKeySpec` object in order to store them. The WebSphere Application Server runtime stores and retrieves the key pair as secret keys. The runtime converts the key pair back to `PublicKey` and `PrivateKey` objects when the server retrieves the pair during the handshake.

Use the following constructors to develop the `com.ibm.websphere.crypto.KeyPair` interface:

- Public and private constructor


```
public KeyPair(java.security.Key publicKey, java.security.Key privateKey)
```
- Certificate and private constructor.


```
public KeyPair(java.security.cert.Certificate[] certChain,
java.security.Key privateKey)
```

The previous example code shows the `KeyPairGenerator` class using the public and private constructor. Each call to this class generates a new and unique key pair, and this class is invoked by a `KeySet` to create a new key pair when `isKeyPair=true`. The version number in the key set increments each time it is called.

Key set groups collection

Use this page to manage groups of public, private, and shared keys. These key groups enable the application server to control multiple sets of Lightweight Third Party Authentication (LTPA) keys.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key set groups**.

Button	Resulting action
New	Adds a key set group. A key set group combines one or more key sets together as a single key set group. It allows the generation of multiple different types of keys to occur at the same time. A single key set represents one type of key, so a key set group allows you to group the different types.
Delete	Deletes an existing key set group. You must be sure that there are no other references to this key set group before you delete it.
Generate keys	Generates keys for key set group. The system generates keys for each key set within the key set group so that the keys remain synchronized with each other in terms of version. You must configure a valid key generation class and a key store that is writable. See the <code>com.ibm.websphere.crypto.KeySetHelper</code> application programming interfaces (APIs) to enable the use of keys that are managed by a <code>KeySetGroup</code> or <code>KeySet</code> .

Key set group name

Specifies the name of the key set group used to reference it.

Automatically generate keys

Specifies that the keys are to be generated automatically on a schedule.

Key set groups settings

Use this page to create new key set groups.

To view this administrative console page, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > {Inbound | Outbound} > ssl_configuration**. Under Related items, click **Key set groups**. Then click the **New** button.

Key set group name

Specifies the name of key set group used. This name can be referenced using the `com.ibm.websphere.crypto.KeySetHelper` API to retrieve the managed keys from an application.

Data type: Text

Key sets

Specifies a set of key instances of the same type for use in cryptographic operations.

Add

Specifies to add the selected key set part of this key set group.

Remove

Specifies to remove the selection from the **Key sets** list.

Automatically generate keys

Specifies that the keys are generated automatically on a schedule. When a new key is generated, the `security.xml` is updated and saved by the runtime to track the key reference version. This can cause save conflicts when updating the same file from admin applications.

Default: Enabled

Scheduled time for generation

Specifies the scheduled time when the system generates selected key set group or groups. You can specify the scheduled time in hours and minutes; specify either A.M. or P.M., or specify 24-hour. You can also specify the day of the week you want the scheduled event to occur. It is recommended that you set this event to occur during a low peak time, especially for keys that are used by runtime for token validation.

Data type: Integer
Default: 0, 0
Range: 1–12, 0–59

Generate on a specific day

Specifies whether to have the generation occur on a specific day of the week. It is best to auto-generate keys during a low peak day.

Default: Enabled

Weekday

Specifies the day of the week on which the expiration monitor will run if the Check on a specific day option is selected.

Default: Sunday
Range: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday

Repeat interval

Specifies the period of time between each schedule time to check for expired certificates or the interval between schedule checks.

Default: Daily
Range: Daily, Weekly

Generate at an interval

Specifies to generate keys at the specified frequency regardless of the day of the week on which generation occurs.

Default: Disabled

Next start date

Specifies the date for the next scheduled check. This allows the deployment manager to be stopped and restarted without resetting the date.

Configuring the Web server plug-in for Secure Sockets Layer

This topic documents the configuration that is necessary to instantiate a secure connection between the Web server plug-in and the internal HTTP transport in the Web container for the Application Server.

Before you begin

WebSphere Application Server has an internal HTTP transport that accepts HTTP requests. If you install an external HTTP server, the Web server plug-in must forward requests from the external HTTP server to Application Server internal HTTP transport. Follow instructions provided by your HTTP vendor to install and configure your HTTP server. Test your HTTP server by accessing `http://your-host-URL` and `https://your-host-URL`. You should also have a Web server plug-in installed. For instructions on installing HTTP Server and Web server plug-in, see *Installing IBM HTTP Server*. If you are installing the HTTP Server and Web server plug-in on Solaris x64, see *Installing IBM HTTP Server* for additional information on enabling the plug-in to load the correct libraries for Secure Sockets Layer (SSL).

1. Create a directory on the Web server host for storing the key ring file that is referenced by the plug-in and associated files, for example: `plugin_install_root/etc/keys`.
2. From the administrative console, click **Servers > Web servers**.
3. Select the Web server name.
4. Click **Plug-in properties**.
5. Click **Manage keys and certificates** to access configuration options for your keys and certificates. By default, you can change your password used to protect the key store.
6. Click **OK**.
7. Click the Web server **keystores** button to copy the key store and to stash files to a managed web server. For non-managed Web servers, use FTP to copy them.
8. Under Additional Properties, you can also select one of the following:
 - **Signer certificates** - Use to add new certificates, delete certificates, extract certificates, and to retrieve certificates from a port.
 - **Personal certificates** - Use to create a new chained or self-signed certificate, delete a certificate, or to import and export a personal certificate.
 - **Personal certificate requests** - Use to manage personal certificate requests.
 - **Custom properties** - Use to define custom properties for the key store.

Results

The IBM HTTP Server plug-in and the internal Web server are configured for SSL.

Web server plug-in default configuration

When you create a new Web server definition, WebSphere Application Server associates the Web server plug-in with a Certificate Management Services (CMS) keystore for a specific node. The keystore contains all of the signers for the current cell with the self-signed or chained certificate, which belongs to the node. The plug-in can communicate securely to WebSphere Application Server, even when the plug-in is configured with Secure Sockets Layer (SSL) client authentication enabled.

When you set the Web server definition to `webserver1` on node `myhostNode01`, WebSphere Application Server creates the keystore configuration. The keystore is scoped to the `webserver1` server, which makes it visible to this server only. Other processes cannot use this keystore definition.

The following sample code from the `security.xml` file shows the configuration entries for the Web server plug-in.

```
<keyStores xmi:id="KeyStore_1132357815719" name="CMSKeyStore"
password="{xor}HRYNFAtRbxEw0zpvbhw6MzM=" provider="IBMCMSProvider"
location="C:\WASX_e0540.11\AppServer\profiles\AppSrv01/config/cells
/myhostCell01/nodes/myhostNode01/servers/webserver1/plugin-key.kdb"
type="CMSKS" fileBased="true" createStashFileForCMS="true"
managementScope="ManagementScope_1132357815718"/>
<managementScopes xmi:id="ManagementScope_1132357815718" scopeName="
(cell):myhostCell01:(node):myhostNode01:(server):webserver1" scopeType="server"/>
```

The following sample code shows how the CMS keystore and stash file are generated in the `security.xml` file.

```
C:\WebSphere\AppServer\profiles\Dmgr01\config\cells\myhostCell01\nodes
\myhostNode01\servers\webserver1\plugin-key.kdb
C:\WebSphere\AppServer\profiles\Dmgr01\config\cells\myhostCell01
\nodes\myhostNode01\servers\webserver1\plugin-key.sth
```

The default password for the keystore is `WebAS`. You can change the default keystore password by using either the administrative console or the appropriate **AdminTask** command. The following sample code shows the **AdminTask** command that you can use to create this CMS keystore.

```
$AdminTask createCMSKeyStore /config/cells/myhostCell01/nodes/myhostNode01
/servers/webserver1/plugin-key.kdb myhost.austin.ibm.com
```

Note the following characteristics of the previous example:

- You can create only one `CMSKeyStore` entry for each management scope. If a CMS keystore already exists for scope `(cell):myhostCell01:(node):myhostNode01:(server):webserver1`, then you cannot create another `CMSKeyStore` entry
- The Uniform Resource Identifier (URI) for the keystore name is `/config/cells/myhostCell01/nodes/myhostNode01/servers/webserver1/plugin-key.kdb`
- The host name in the plug-in location is `myhost.austin.ibm.com`. WebSphere Application Server uses this name to create a chained certificate, if a chained certificate does not already exist for that particular node. If a chained certificate already exists for the node, then the certificate is put into the CMS keystore and all the signers from the cell are added, by default.

When additional nodes are federated, the signers for these nodes are not automatically added to each Web server for the CMS keystore. For the Web server plug-in to be able to communicate with a newly federated node, you must manually exchange signers with the `CMSKeyStore` keystore. Use the administrative console keystore certificate management function to exchange signers. For more information, see “Extracting a signer certificate from a personal certificate” on page 700.

Chapter 8. Developing extensions to the WebSphere security infrastructure

WebSphere Application Server provides various plug points so that you can extend the security infrastructure. Extending this security infrastructure involves several activities including: Developing custom user registries, developing applications that use programmatic security, and customizing Web application login forms.

About this task

The following topics are covered in this section:

- Developing custom user registries
- Developing applications that use programmatic security
- Customizing Web application login forms
- Customizing application login forms with Java Authentication and Authorization Service (JAAS)
- Securing transports with Java Secure Sockets Extension (JSSE) and Java Cryptography Extension (JCE) programming interfaces
- Implementing tokens for security attribute propagation

Developing standalone custom registries

This development provides considerable flexibility in adapting WebSphere Application Server security to various environments where some notion of a user registry, other than LDAP or Local OS, already exists in the operational environment.

Before you begin

WebSphere Application Server security supports the use of standalone custom registries in addition to the local operating system registry, standalone Lightweight Directory Access Protocol (LDAP) registries, and federated repositories for authentication and authorization purposes. A standalone custom-implemented registry uses the `UserRegistry` Java interface as provided by WebSphere Application Server. A standalone custom-implemented registry can support virtually any type or notion of an accounts repository from a relational database, flat file, and so on.

Implementing a standalone custom registry is a software development effort. Implement the methods that are defined in the `com.ibm.websphere.security.UserRegistry` interface to make calls to the appropriate registry to obtain user and group information. The interface defines a general set of methods for encapsulating a wide variety of registries. You can configure a standalone custom registry as the selected repository when configuring WebSphere Application Server security on the Global security panel.

In WebSphere Application Server Version 7.0, make sure that your implementation of the standalone custom registry does not depend on any WebSphere Application Server components such as data sources, Enterprise JavaBeans (EJB) and Java Naming and Directory Interface (JNDI). You can not have this dependency because security is initialized and enabled prior to most of the other WebSphere Application Server components during startup. If your previous implementation used these components, make a change that eliminates the dependency. For example, if your previous implementation used data sources to connect to a database, instead use the `JDBC java.sql.DriverManager` interface to connect to the database.

Refer to the “Migrating custom user registries” on page 20 for more information on migrating. If your previous implementation uses data sources to connect to a database, change the implementation to use Java database connectivity (JDBC) connections.

1. Implement all the methods in the interface except for the `CreateCredential` method, which is implemented by WebSphere Application Server. “FileRegistrySample.java file” on page 147 is provided for reference.

Note: The sample provided is intended to familiarize you with this feature. Do not use this sample in an actual production environment.

2. Build your implementation.

To compile your code, you need the `app_server_install_root/Base/plugins/com.ibm.ws.runtime.jar` and the `app_server_install_root/Base/plugins/com.ibm.ws.security.crypto.jar` files in your class path. For example:

```
%install_root%/java/bin/javac -classpath
app_server_install_root/Base/plugins/com.ibm.ws.runtime.jar;
app_server_install_root/Base/plugins/com.ibm.ws.security.crypto.jar your_implementation_file.java
```

3. Copy the class files that are generated in the previous step to the product class path.

The preferred location is the `%install_root%/lib/ext` directory. Copy these class files to all of the product process class paths.

4. Follow the steps in “Configuring standalone custom registries” on page 142 to configure your implementation using the administrative console. This step is required to implement custom user registries.

What to do next

If you enable security, make sure that you complete the remaining steps:

1. Save and synchronize the configuration and restart all of the servers.
2. Try accessing some J2EE resources to verify that the custom registry implementation is correct.

Example: Standalone custom registries

Use these links to view registry examples.

A *standalone custom registry* is a customer-implemented registry that implements the `UserRegistry` Java interface, as provided by WebSphere Application Server. A custom-implemented registry can support virtually any type or form of an accounts repository from a relational database, flat file, and so on. The custom registry provides considerable flexibility in adapting WebSphere Application Server security to various environments where some form of a registry, other than a federated repository, Lightweight Directory Access Protocol (LDAP) registry, or local operating system registry, already exist in the operational environment.

To view a sample standalone custom registry, refer to the following files:

- “FileRegistrySample.java file” on page 147
- “users.props file” on page 164
- “groups.props file” on page 164

Result.java file

This module is used by user registries in WebSphere Application Server when calling the `getUsers` and `getGroups` methods. The user registries use this method to set the list of users and groups and to indicate if more users and groups in the user registry exist than requested.

```
//
// 5639-D57, 5630-A36, 5630-A37, 5724-D18
// (C) COPYRIGHT International Business Machines Corp. 1997, 2005
// All Rights Reserved * Licensed Materials - Property of IBM
//
package com.ibm.websphere.security;

import java.util.List;
```



```

public class Result implements java.io.Serializable {
    /**
     * Default constructor
     */
    public Result() {
    }

    /**
     * Returns the list of users and groups
     * @return the list of users and groups
     */
    public List getList() {
        return list;
    }

    /**
     * indicates if there are more users and groups in the registry
     */
    public boolean hasMore() {
        return more;
    }

    /**
     * Set the flag to indicate that there are more users and groups
     * in the registry to true
     */
    public void setHasMore() {
        more = true;
    }

    /**
     * Set the list of users and groups
     * @param list list of users/groups
     */
    public void setList(List list) {
        this.list = list;
    }

    private boolean more = false;
    private List list;
}

```

UserRegistry.java files

The following file is a custom property that is used with a custom user registry.

For more information, see “Configuring standalone custom registries” on page 142.

```

// 5639-D57, 5630-A36, 5630-A37, 5724-D18
// (C) COPYRIGHT International Business Machines Corp. 1997, 2005
// All Rights Reserved * Licensed Materials - Property of IBM
//
// DESCRIPTION:
//
// This file is the UserRegistry interface that custom registries in WebSphere
// Application Server implement to enable WebSphere security to use the custom
// registry.
//
package com.ibm.websphere.security;

import java.util.*;
import java.rmi.*;
import java.security.cert.X509Certificate;
import com.ibm.websphere.security.cred.WSCredential;/**
 * Implementing this interface enables WebSphere Application Server Security
 * to use custom registries. This interface extends java.rmi.Remote because the
 * registry can be in a remote process.

```

```

*
* Implementation of this interface must provide implementations for:
*
* initialize(java.util.Properties)
* checkPassword(String,String)
* mapCertificate(X509Certificate[])
* getRealm
* getUsers(String,int)
* getUserDisplayName(String)
* getUniqueId(String)
* getUserSecurityName(String)
* isValidUser(String)
* getGroups(String,int)
* getGroupDisplayName(String)
* getUniqueGroupId(String)
* getUniqueGroupIds(String)
* getGroupSecurityName(String)
* isValidGroup(String)
* getGroupsForUser(String)
* getUsersForGroup(String,int)
* createCredential(String)
**/

public interface UserRegistry extends java.rmi.Remote
{

    /**
     * Initializes the registry. This method is called when creating the
     * registry.
     *
     * @param props the registry-specific properties with which to
     *             initialize the custom registry
     * @exception CustomRegistryException
     *             if there is any registry specific problem
     * @exception RemoteException
     *             as this extends java.rmi.Remote
     **/
    public void initialize(java.util.Properties props)
        throws CustomRegistryException,
            RemoteException; /**
     * Checks the password of the user. This method is called to authenticate a
     * user when the user's name and password are given.
     *
     * @param userSecurityName the name of the user
     * @param password the password of the user
     * @return a valid userSecurityName. Normally this is
     *         the name of same user whose password was checked but if the
     *         implementation wants to return any other valid
     *         userSecurityName in the registry it can do so
     * @exception CheckPasswordFailedException if userSecurityName/
     *         password combination does not exist in the registry
     * @exception CustomRegistryException if there is any registry specific
     *         problem
     * @exception RemoteException as this extends java.rmi.Remote
     **/
    public String checkPassword(String userSecurityName, String password)
        throws PasswordCheckFailedException,
            CustomRegistryException,
            RemoteException; /**
     * Maps a certificate (of X509 format) to a valid user in the registry.
     * This is used to map the name in the certificate supplied by a browser
     * to a valid userSecurityName in the registry
     *
     */

```

```

* @param cert the X509 certificate chain
* @return the mapped name of the user userSecurityName
* @exception CertificateMapNotSupportedException if the particular
*         certificate is not supported.
* @exception CertificateMapFailedException if the mapping of the
*         certificate fails.
* @exception CustomRegistryException if there is any registry specific
*         problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String mapCertificate(X509Certificate[] cert)
    throws CertificateMapNotSupportedException,
           CertificateMapFailedException,
           CustomRegistryException,
           RemoteException; /**
* Returns the realm of the registry.
*
* @return the realm. The realm is a registry-specific string indicating
*         the realm or domain for which this registry
*         applies. For example, for OS400 or AIX this would be the
*         host name of the system whose user registry this object
*         represents.
*         If null is returned by this method realm defaults to the
*         value of "customRealm". It is recommended that you use
*         your own value for realm.
* @exception CustomRegistryException if there is any registry specific
*         problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getRealm()
    throws CustomRegistryException,
           RemoteException; /**
* Gets a list of users that match a pattern in the registry.
* The maximum number of users returned is defined by the limit
* argument.
* This method is called by administrative console and by scripting (command
* line) to make available the users in the registry for adding them (users)
* to roles.
*
* @parameter pattern the pattern to match. (For example., a* will match all
*         userSecurityNames starting with a)
* @parameter limit the maximum number of users that should be returned.
*         This is very useful in situations where there are thousands of
*         users in the registry and getting all of them at once is not
*         practical. A value of 0 implies get all the users and hence
*         must be used with care.
* @return a Result object that contains the list of users
*         requested and a flag to indicate if more users exist.
* @exception CustomRegistryException if there is any registry specific
*         problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public Result getUsers(String pattern, int limit)
    throws CustomRegistryException,
           RemoteException; /**
* Returns the display name for the user specified by userSecurityName.
*
* This method is called only when the user information displays
* (information purposes only, for example, in the administrative console) and not used
* in the actual authentication or authorization purposes. If there are no
* display names in the registry return null or empty string.
*
* In WebSphere Application Server Version 4.0 custom registry, if you had a display

```

```

* name for the user and if it was different from the security name, the display name
* was returned for the EJB methods getCallerPrincipal() and the servlet methods
* getUserPrincipal() and getRemoteUser().
* In WebSphere Application Server Version 5.0 for the same methods the security
* name is returned by default. This is the recommended way as the display name
* is not unique and might create security holes.
*
* See the documentation for more information.
*
* @parameter userSecurityName the name of the user.
* @return the display name for the user. The display name
* is a registry-specific string that represents a descriptive, not
* necessarily unique, name for a user. If a display name does
* not exist return null or empty string.
* @exception EntryNotFoundException if userSecurityName does not exist.
* @exception CustomRegistryException if there is any registry specific
* problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getUserDisplayName(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException; /**
* Returns the unique ID for a userSecurityName. This method is called when
* creating a credential for a user.
*
* @parameter userSecurityName the name of the user.
* @return the unique ID of the user. The unique ID for a user is
* the stringified form of some unique, registry-specific, data
* that serves to represent the user. For example, for the UNIX
* user registry, the unique ID for a user can be the UID.
* @exception EntryNotFoundException if userSecurityName does not exist.
* @exception CustomRegistryException if there is any registry specific
* problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getUniqueUserId(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException; /**
* Returns the name for a user given its unique ID.
*
* @parameter uniqueUserId the unique ID of the user.
* @return the userSecurityName of the user.
* @exception EntryNotFoundException if the uniqueUserID does not exist.
* @exception CustomRegistryException if there is any registry specific
* problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getUserSecurityName(String uniqueUserId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
* Determines if the userSecurityName exists in the registry
*
* @parameter userSecurityName the name of the user
* @return true if the user is valid. false otherwise
* @exception CustomRegistryException if there is any registry specific
* problem
* @exception RemoteException as this extends java.rmi.Remote
**/

```

```

public boolean isValidUser(String userSecurityName)
    throws CustomRegistryException,
           RemoteException;

/**
 * Gets a list of groups that match a pattern in the registry.
 * The maximum number of groups returned is defined by the limit
 * argument.
 * This method is called by the administrative console and scripting
 * (command line) to make available the groups in the registry for adding
 * them (groups) to roles.
 *
 * @parameter pattern the pattern to match. (For e.g., a* will match all
 * groupSecurityNames starting with a)
 * @parameter limit the maximum number of groups to return.
 * This is very useful in situations where there are thousands of
 * groups in the registry and getting all of them at once is not
 * practical. A value of 0 implies get all the groups and hence
 * must be used with care.
 * @return a Result object that contains the list of groups
 * requested and a flag to indicate if more groups exist.
 * @exception CustomRegistryException if there is any registry-specific
 * problem
 * @exception RemoteException as this extends java.rmi.Remote
 */
public Result getGroups(String pattern, int limit)
    throws CustomRegistryException,
           RemoteException;

/**
 * Returns the display name for the group specified by groupSecurityName.
 *
 * This method may be called only when the group information displayed
 * (for example, the administrative console) and not used in the actual
 * authentication or authorization purposes. If there are no display names
 * in the registry return null or empty string.
 *
 * @parameter groupSecurityName the name of the group.
 * @return the display name for the group. The display name
 * is a registry-specific string that represents a descriptive, not
 * necessarily unique, name for a group. If a display name does
 * not exist return null or empty string.
 * @exception EntryNotFoundException if groupSecurityName does not exist.
 * @exception CustomRegistryException if there is any registry specific
 * problem
 * @exception RemoteException as this extends java.rmi.Remote
 */
public String getGroupDisplayName(String groupSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
 * Returns the unique ID for a group.
 *
 * @parameter groupSecurityName the name of the group.
 * @return the unique ID of the group. The unique ID for
 * a group is the stringified form of some unique,
 * registry-specific, data that serves to represent the group.
 * For example, for the UNIX user registry, the unique ID might
 * be the GID.
 * @exception EntryNotFoundException if groupSecurityName does not exist.
 * @exception CustomRegistryException if there is any registry specific

```

```

*           problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getUniqueGroupId(String groupSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
* Returns the unique IDs for all the groups that contain the unique ID of
* a user.
* Called during creation of a user's credential.
*
* @parameter uniqueUserId the unique ID of the user.
* @return a list of all the group unique IDs that the unique user ID
* belongs to. The unique ID for an entry is the stringified
* form of some unique, registry-specific, data that serves
* to represent the entry. For example, for the
* UNIX user registry, the unique ID for a group could be the GID
* and the unique ID for the user could be the UID.
* @exception EntryNotFoundException if unique user ID does not exist.
* @exception CustomRegistryException if there is any registry specific
*           problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public List getUniqueGroupIds(String uniqueUserId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
* Returns the name for a group given its unique ID.
*
* @parameter uniqueGroupId the unique ID of the group.
* @return the name of the group.
* @exception EntryNotFoundException if the uniqueGroupId does not exist.
* @exception CustomRegistryException if there is any registry-specific
*           problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public String getGroupSecurityName(String uniqueGroupId)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
* Determines if the groupSecurityName exists in the registry
*
* @parameter groupSecurityName the name of the group
* @return true if the groups exists, false otherwise
* @exception CustomRegistryException if there is any registry specific
*           problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public boolean isValidGroup(String groupSecurityName)
    throws CustomRegistryException,
           RemoteException;

/**
* Returns the securityNames of all the groups that contain the user
*
* This method is called by administrative console and scripting

```

```

* (command line) to verify the user entered for RunAsRole mapping belongs
* to that role in the roles to user mapping. Initially, the check is done
* to see if the role contains the user. If the role does not contain the user
* explicitly, this method is called to get the groups that this user
* belongs to so that checks are made on the groups that the role contains.
*
* @parameter userSecurityName the name of the user
* @return a List of all the group securityNames that the user
* belongs to.
* @exception EntryNotFoundException if user does not exist.
* @exception CustomRegistryException if there is any registry specific
* problem
* @exception RemoteException as this extends java.rmi.Remote
**/
public List getGroupsForUser(String userSecurityName)
    throws EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

/**
* Gets a list of users in a group.
*
* The maximum number of users returned is defined by the limit
* argument.
*
* This method is used by the WebSphere Business Integration
* Server Foundation process choreographer when staff assignments
* are modeled using groups.
*
* In rare situations where you are working with a user registry and it is not
* practical to get all of the users from any of your groups (for example if
* a large number of users exist) you can create the NotImplementedException
* for those particular groups. Make sure that if the WebSphere Business
* Integration Server Foundation Process Choreographer is installed (or
* if installed later) that the users are not modeled using these particular groups.
* If no concern exists about the staff assignments returning the users from
* groups in the registry it is recommended that this method be implemented
* without throwing the NotImplementedException.
*
* @parameter groupSecurityName that represents the name of the group
* @parameter limit the maximum number of users to return.
* This option is very useful in situations where lots of
* users are in the registry and getting all of them at
* once is not practical. A value of 0 means get all of
* the users and must be used with care.
* @return a Result object that contains the list of users
* requested and a flag to indicate if more users exist.
* @deprecated This method will be deprecated in the future.
* @exception NotImplementedException create this exception in rare situations
* if it is not practical to get this information for any of the
* groups from the registry.
* @exception EntryNotFoundException if the group does not exist in
* the registry
* @exception CustomRegistryException if any registry-specific
* problem occurs
* @exception RemoteException as this extends java.rmi.Remote interface
**/
public Result getUsersForGroup(String groupSecurityName, int limit)
    throws NotImplementedException,
           EntryNotFoundException,
           CustomRegistryException,
           RemoteException;

```



```

/**
 * This method is implemented internally by the WebSphere Application Server
 * code in this release. This method is not called for the custom registry
 * implementations for this release. Return null in the implementation.
 *
 * Note that because this method is not called you can also return the
 * NotImplementedException as the previous documentation says.
 *
 */
public com.ibm.websphere.security.cred.WSCredential
        createCredential(String userSecurityName)
        throws NotImplementedException,
        EntryNotFoundException,
        CustomRegistryException,
        RemoteException;
}

```

Implementing custom password encryption

WebSphere Application Server supports the use of custom password encryption.

Before you begin

An installation can implement any password encryption algorithm it chooses.

About this task

Complete the following steps to implement custom password encryption:

1. Build your custom password encryption class. An example of a custom password encryption class follows.

```

// CustomPasswordEncryption
// Encryption and decryption functions
public interface CustomPasswordEncryption {
    public EncryptedInfo encrypt(byte[] clearText) throws PasswordEncryptException;
    public byte[] decrypt(EncryptedInfo cipherTextInfo) throws PasswordEncryptException;
    public void initialize(HashMap initParameters);
};
// Encapsulation of cipher text and label
public class EncryptedInfo {
    public EncryptedInfo(byte[] bytes, String keyAlias);
    public byte[] getEncryptedBytes();
    public String getKeyAlias();
};

```

2. Enable custom password encryption.
 - a. Set the custom property **com.ibm.wsspi.security.crypto.customPasswordEncryptionClass** to the name of the class that is to be given control.
 - b. Enable the function. Set the custom property, **com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled** to true.

Results

Custom password encryption at the installation is complete.

Developing applications that use programmatic security

For some applications, declarative security is not sufficient to express the security model of the application. Use this topic to develop applications that use programmatic security.

About this task

IBM WebSphere Application Server provides security components that provide or collaborate with other services to provide authentication, authorization, delegation, and data protection. WebSphere Application Server also supports the security features that are described in the Java Platform, Enterprise Edition (Java EE) specification. An application goes through three stages before it is ready to run:

- Development
- Assembly
- Deployment

Most of the security for an application is configured during the assembly stage. The security that is configured during the assembly stage is called *declarative security* because the security is *declared* or *defined* in the deployment descriptors. The declarative security is enforced by the security runtime. For some applications, declarative security is not sufficient to express the security model of the application. For these applications, you can use *programmatic security*.

1. Develop secure Web applications. For more information, see “Developing with programmatic security APIs for Web applications” on page 764.
2. Develop servlet filters for form login processing. For more information, see “Developing servlet filters for form login processing” on page 778.
3. Develop form login pages. For more information, see “Customizing Web application login” on page 773.
4. Develop enterprise bean component applications. For more information, see “Developing with programmatic APIs for EJB applications” on page 769.
5. Develop with Java Authentication and Authorization Service to log in programmatically. For more information, see “Developing programmatic logins with the Java Authentication and Authorization Service” on page 389.
6. Develop your own Java EE security mapping module. For more information, see “Configuring programmatic logins for Java Authentication and Authorization Service” on page 403.
7. Develop custom user registries. For more information, see “Developing standalone custom registries” on page 731.
8. Develop a custom interceptor for trust associations.

Protecting system resources and APIs (Java 2 security)

Java 2 security is a programming model that is very pervasive and has a huge impact on application development.

Before you begin

Java 2 security is orthogonal to Java Platform, Enterprise Edition (Java EE) role-based security; you can disable or enable it independently of administrative security.

However, it does provide an extra level of access control protection on top of the Java EE role-based authorization. It particularly addresses the protection of system resources and application programming interfaces (API). Administrators need to consider the benefits against the risks of disabling Java 2 security.

The following recommendations are provided to help enable Java 2 security in a test or production environment:

1. Make sure the application is developed with the Java 2 security programming model. Developers have to know whether or not the APIs that are used in the applications are protected by Java 2 security. It is very important that the required permissions for the APIs used are declared in the policy file (`was.policy`), or the application fails to run when Java 2 security is enabled. Developers can reference the Web site for Development Kit APIs that are protected by Java 2 security. See the Programming model and decisions section of the Security: Resources for learning topic to visit this Web site.

2. Make sure that migrated applications from previous releases are given the required permissions. Because Java 2 security is not supported or partially supported in previous WebSphere Application Server releases, applications developed prior to Version 5 most likely are not using the Java 2 security programming model. No easy way to find out all the required permissions for the application is available. The following are activities you can perform to determine the extra permissions that are required by an application:
 - Code review and code inspection
 - Application documentation review
 - Sandbox testing of migrated enterprise applications with Java 2 security enabled in a preproduction environment. Enable tracing in WebSphere Java 2 security manager to help determine the missing permissions in the application policy file. The trace specification is:
`com.ibm.ws.security.core.SecurityManager=all=enabled.`
 - Use the `com.ibm.websphere.java2secman.norethrow` system property to aid debugging. Do not use this property in a production environment.Refer to “Java 2 security” on page 44

The default permission set for applications is the recommended permission set that is defined in the J2EE 1.3 Specification. The default is declared in the `app_server_root/profiles/profile_name/config/cells/cell_name/nodes/node_name/app.policy` policy file with permissions defined in the Development Kit (`JAVA_HOME/jre/lib/security/java.policy`) policy file that grant permissions to everyone. However, applications are denied permissions that are declared in the `profiles/profile_name/config/cells/cell_name/filter.policy` file. Permissions that are declared in the `filter.policy` file are filtered for applications during the permission check.

Define the required permissions for an application in a `was.policy` file and embed the `was.policy` file in the application enterprise archive (EAR) file as `YOURAPP.ear/META-INF/was.policy`, see “Configuring Java 2 security policy files” on page 744 for details.

The following steps describe how to enforce Java 2 security on the cell level for WebSphere Application Server Network Deployment and the server level for WebSphere Application Server and WebSphere Application Server Express:

1. Click **Security > Global security**. The Global security panel is displayed.
2. Select the **Use Java 2 security to restrict application access to local resources** option.
3. Click **OK** or **Apply**.
4. Click **Save** to save the changes.
5. Restart the server for the changes to take effect.

Results

Java 2 security is enabled and enforced for the servers. Java 2 security permission is selected when a Java 2 security protected API is called.

When to use Java 2 security

1. Enable protection on system resources, for example when opening or listening to a socket connection, reading or writing to operating system file systems, reading or writing Java virtual machine system properties, and so on.
2. Prevent application code from calling destructive APIs, for example, calling the `System.exit` method brings down the application server.
3. Prevent application code from obtaining privileged information (passwords) or gaining extra privileges (obtaining server credentials).

What to do next

You can enforce Java 2 security on the server level for WebSphere Application Server Network Deployment by completing the following steps.

Note: Changes to Java 2 security settings on the server level override the settings on the cell level.

1. Click **Servers > Application servers > *server_name***.
2. Under Security, click **Server security**.
3. Select the **Security settings for this server override cell settings** option.
4. Select the **Use Java 2 security to restrict application access to local resources** option.
5. Click **OK** or **Apply**.
6. Click **Save** to save the changes.
7. Restart the server for the changes to take effect.

The Java 2 security manager is enhanced to dump the Java 2 security permissions that are granted to all classes on the call stack when an application is denied access to a resource. The `java.security.AccessControlException` exception is created. However, this tracing capability is disabled by default. You can enable this capability by specifying the server trace service with the `com.ibm.ws.security.core.SecurityManager=all=enabled` trace specification. When the exception is created, the trace dump provides hints to determine whether the application is missing permissions or the product runtime code or the third-party libraries that are used are not properly marked as `privileged` when accessing Java 2 security-protected resources. See the Security Problem Determination Guide for details.

Using PolicyTool to edit policy files

Use the **PolicyTool** utility to update policy files.

Before you begin

Java 2 security uses several policy files to determine the granted permission for each Java program. The Java Development Kit provides the **PolicyTool** tool to edit these policy files. This tool is recommended for editing any policy file to verify the syntax of its contents. Syntax errors in the policy file cause an `AccessControlException` exception when the application runs, including the server start. Identifying the cause of this exception is not easy because the user might not be familiar with the resource that has an access violation. Be careful when you edit these policy files.

See “Java 2 security policy files” on page 48 for the list of available policy files.

1. Start the **PolicyTool**.

Windows For example, you can enter the following command at a Windows command prompt:

```
%{was.install.root}/java/jre/bin/policytool
```

The **PolicyTool** window opens. The tool looks for the `java.policy` file in your home directory. If it does not exist, an error message displays.

Click **OK**.

2. Click **File > Open**.
3. Navigate the directory tree in the **Open** window to pick up the policy file that you need to update. After selecting the policy file, click **Open**. The code base entries are listed in the window.
4. Create or modify the code base entry.
 - a. Modify the existing code base entry by double-clicking the code base, or click the code base and click **Edit Policy Entry**. The Policy Entry window opens with the permission list defined for the selected code base.
 - b. Create a new code base entry by clicking **Add Policy Entry**.

The Policy Entry window opens. At the code base column, enter the code base information as a URL format.

For example, you can enter:

```
app_server_root/InstalledApps/testcase.ear
```

where the *app_server_root* variable represents your installation location.

5. Modify or add the permission specification.
 - a. Modify the permission specification by double-clicking the entry that you want to modify, or by selecting the permission and clicking **Edit Permission**. The Permissions window opens with the selected permission information.
 - b. Add a new permission by clicking **Add Permission**. The Permissions window opens. In the Permissions window are four rows for Permission, Target Name, Actions, and Signed By.
6. Select the permission from the Permission list. The selected permission displays. After a permission is selected, the Target Name, Actions, and Signed By fields automatically show the valid choices or they enable text input in the right text input area.
 - a. Select **Target Name** from the list, or enter the target name in the right text input area.
 - b. Select **Actions** from the list.
 - c. Input **Signed By** if it is needed.

Note: The Signed By keyword is not supported in the following policy files: *app.policy*, *spi.policy*, *library.policy*, *was.policy*, and *filter.policy* files. However, the Signed By keyword is supported in the following policy files: *#java.policy*, *server.policy*, and *client.policy* files. The Java Authentication and Authorization Service (JAAS) is not supported in the *app.policy*, *spi.policy*, *library.policy*, *was.policy*, and *filter.policy* files. However, the JAAS principal keyword is supported in a JAAS policy file when it is specified by the *java.security.auth.policy* Java virtual machine (JVM) system property.

7. Click **OK** to close the Permissions window. Modified permission entries of the specified code base display.
8. Click **Done** to close the window. Modified code base entries are listed. Repeat the previous steps until you complete editing.
9. Click **File** > **Save** after you finish editing the file.

Results

A policy file is updated. If any policy files need editing, use the **PolicyTool** utility. Do not edit the policy file manually. Syntax errors in the policy files can potentially cause application servers or enterprise applications to not start or function incorrectly. For the changes in the updated policy file to take effect, restart the Java processes.

Configuring Java 2 security policy files

Users can configure Java 2 security policy files so that the required permission is granted for the specified WebSphere Application Server enterprise application.

Before you begin

Java 2 security uses several policy files to determine the permissions for each Java programs.

Two types of policy files are supported by WebSphere Application Server: dynamic policy files and static policy files. Static policy files provide the default permissions. Dynamic policy files provide application permissions. Six dynamic policy files are provided:

Policy file name	Description
app.policy	Contains default permissions for all of the enterprise applications in the cell. Note: Updates to the app.policy file only apply to the enterprise applications on the node to which the app.policy file belongs.
was.policy	Contains application-specific permissions for an WebSphere Application Server enterprise application. This file is packaged in an enterprise archive (EAR) file.
ra.xml	Contains connector application specific permissions for a WebSphere Application Server enterprise application. This file is packaged in a resource adapter archive (RAR) file.
spi.policy	Contains permissions for Service Provider Interface (SPI) or third-party resources that are embedded in WebSphere Application Server. The default contents grant everything. Update this file carefully when the cell requires more protection against SPI in the cell. This file is applied to all of the SPIs that are defined in the resources.xml file.
library.policy	Contains permissions for the shared library of enterprise applications.
filter.policy	Contains the list of permissions that require filtering from the was.policy file and the app.policy file in the cell. This filtering mechanism only applies to the was.policy and app.policy files.

In WebSphere Application Server, applications must have the appropriate thread permissions specified in the was.policy or app.policy file. Without the thread permissions specified, the application cannot manipulate threads and WebSphere Application Server creates a java.security.AccessControlException exception. The app.policy file applies to a specified node. If you change the permissions in one app.policy file, you must incorporate the new thread policy in the same file on the remaining nodes. Also, if you add the thread permissions to the app.policy file, you must restart WebSphere Application Server to enforce the new permissions. However, if you add the permissions to the was.policy file for a specific application, you do not need to restart WebSphere Application Server. An administrator must add the following code to a was.policy or app.policy file for an application to manipulate threads:

```
grant codeBase "file:${application}" {
    permission java.lang.RuntimePermission "stopThread";
    permission java.lang.RuntimePermission "modifyThread";
    permission java.lang.RuntimePermission "modifyThreadGroup";
};
```

Note: The Signed By keyword is not supported in the following policy files: app.policy, spi.policy, library.policy, was.policy, and filter.policy files. However, the Signed By keyword is supported in the following policy files: java.policy, server.policy, and client.policy files. The Java Authentication and Authorization Service (JAAS) is not supported in the app.policy, spi.policy, library.policy, was.policy, and filter.policy files. However, the JAAS principal keyword is supported in a JAAS policy file when it is specified by the java.security.auth.policy Java virtual machine (JVM) system property. You can statically set the authorization policy files in java.security.auth.policy with auth.policy.url.n=URL, where URL is the location of the authorization policy.

1. Identify the policy file to update.

- If the permission is required by an application, update the static policy file. Refer to “Configuring static policy files” on page 758.
- If the permission is required by all of the WebSphere Application Server enterprise applications in the node, refer to “spi.policy file permissions” on page 753.
- If the permission is required only by specific WebSphere Application Server enterprise applications and the permission is required only by connector, update the ra.xml file. Refer to Assembling resource adapter (connector) modules. Otherwise, update the was.policy file. Refer to “Configuring the was.policy file” on page 750 and “Adding the was.policy file to applications” on page 756.
- If the permission is required by shared libraries, refer to “library.policy file permissions” on page 754.
- If the permission is required by SPI libraries, refer to “spi.policy file permissions” on page 753.

Note: Pick up the policy file with the smallest scope. You can avoid giving an extra permission to the Java programs and protect the resources. You can update the `ra.xml` file or the `was.policy` file rather than the `app.policy` file. Use specific component symbols (`$(ejbcomponent)`, `$(webComponent)`, `$(connectorComponent)` and `$(jars)`) than `$(application)` symbols. Update dynamic policy files, rather than static policy files.

Add any permission that you never want granted to the WebSphere Application Server enterprise application in the cell to the `filter.policy` file. Refer to “filter.policy file permissions” on page 749.

2. Restart the WebSphere Application Server enterprise application.

Results

The required permission is granted for the specified WebSphere Application Server enterprise application.

Example

If an WebSphere Application Server enterprise application in a cell requires permissions, some of the dynamic policy files need updating. The symptom of the missing permission is the `java.security.AccessControlException` exception. The missing permission is listed in the exception data, which will appear as one line, but is split in sections below for readability.

When a Java program receives this exception and adding this permission is justified, add a permission to an adequate dynamic policy file.

The previous permission information lines are split for the illustration. Enter the permission on one line.

To decide whether to add a permission, refer to the “Access control exception” on page 52 topic.

app.policy file permissions:

Java 2 security uses several policy files to determine the granted permissions for each Java program. The union of the permissions that are contained in these following files is applied to the WebSphere Application Server enterprise application. This union determines the granted permissions.

For the list of available policy files that are supported by WebSphere Application Server, see the “Java 2 security policy files” on page 48 article. The `app.policy` file is a default policy file that is shared by all of the WebSphere Application Server enterprise applications. The union of the permissions that are contained in the following files is applied to the WebSphere Application Server enterprise application:

- Any policy file that is specified in the `policy.url.*` properties in the `java.security` file.
- The `app.policy` files, which are managed by configuration and file replication services.
- The `server.policy` file.
- The `java.policy` file.
- The application `was.policy` file.
- The permission specification of the `ra.xml` file.
- The shared library, which is the `library.policy` file.

Changes made in these files are replicated to other nodes in the Network Deployment cell.

In WebSphere Application Server, applications that manipulate threads must have the appropriate thread permissions specified in the `was.policy` or `app.policy` file. Without the thread permissions specified, the application cannot manipulate threads and WebSphere Application Server creates a `java.security.AccessControlException` exception. If an administrator adds thread permissions to the `app.policy` file, the permission change requires a restart of the WebSphere Application Server. An administrator must add the following code to a `was.policy` or `app.policy` file for an application to manipulate threads:


```
grant codeBase "file:${application}" {
    permission java.lang.RuntimePermission "stopThread";
    permission java.lang.RuntimePermission "modifyThread";
    permission java.lang.RuntimePermission "modifyThreadGroup";
};
```

Note: The Signed By and the Java Authentication and Authorization Service (JAAS) principal keywords are not supported in the `app.policy` file. However, the Signed By keyword is supported in the following files: `java.policy`, `server.policy`, and the `client.policy` files. The JAAS principal keyword is supported in a JAAS policy file when it is specified by the `java.security.auth.policy` Java virtual machine (JVM) system property. You can statically set the authorization policy files in the `java.security.auth.policy` property with `auth.policy.url.n=URL` where `URL` is the location of the authorization policy.

If the default permissions for enterprise applications (the union of the permissions that is defined in the `java.policy` file, the `server.policy` file and the `app.policy` file) are enough; no action is required. The default `app.policy` file is used automatically. If a specific change is required to all of the enterprise applications in the cell, update the `app.policy` file. Syntax errors in the policy files cause start failures in the application servers. Edit these policy files carefully.

Note: Updates to the `app.policy` file only apply to the enterprise applications on the node to which the `app.policy` file belongs.

To extract the policy file, use a command prompt to enter the following command on one line using the appropriate variable values for your environment:

Windows

```
wsadmin> set obj [$AdminConfig extract cells/cell_name/node/node_name/app.policy c:\temp\test\app.policy]
```

Edit the extracted `app.policy` file with the Policy Tool. For more information, see “Using PolicyTool to edit policy files” on page 743. Changes to the `app.policy` file are local for the node.

To check in the policy file, use a command prompt to enter the following command on one line using the appropriate variable values for your environment:

Windows

```
wsadmin> $AdminConfig checkin cells/cell_name/nodes/node_name/app.policy c:\temp\test\app.policy $obj
```

Several product-reserved symbols are defined to associate the permission lists to a specific type of resource.

Symbol	Meaning
<code>file:\${application}</code>	Permissions apply to all resources within the application
<code>file:\${jars}</code>	Permissions apply to all utility Java archive (JAR) files within the application
<code>file:\${ejbComponent}</code>	Permissions apply to enterprise bean resources within the application
<code>file:\${webComponent}</code>	Permissions apply to Web resources within the application
<code>file:\${connectorComponent}</code>	Permissions apply to connector resources both within the application and within standalone connector resources.

Five embedded symbols are provided to specify the path and name for the `java.io.FilePermission` permission. These symbols enable flexible permission specifications. The absolute file path is fixed after the installation of the application.

Symbol	Meaning
<code>\${app.installed.path}</code>	Path where the application is installed
<code>\${was.module.path}</code>	Path where the module is installed
<code>\${current.cell.name}</code>	Current cell name
<code>\${current.node.name}</code>	Current node name
<code>\${current.server.name}</code>	Current server name

Note: You cannot use the `${was.module.path}` in the `${application}` entry.

Note: In the following code sample, the first two lines that are related to `java.io.FilePermission` permission are split into two lines for illustrative purposes only.

```
grant codeBase "file:${application}" {
    // The following are required by JavaMail
    permission java.io.FilePermission "${was.install.root}${lib}/activation-impl.jar", "read";
    permission java.io.FilePermission "${was.install.root}${lib}/mail-impl.jar", "read";
};

grant codeBase "file:${jars}" {
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

grant codeBase "file:${connectorComponent}" {
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

grant codeBase "file:${webComponent}" {
    permission java.io.FilePermission "${was.module.path}${/}-", "read, write";
    permission java.lang.RuntimePermission "loadLibrary.*";
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

grant codeBase "file:${ejbComponent}" {
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};
```

If all of the WebSphere Application Server enterprise applications in a cell require permissions that are not defined as defaults in the `java.policy` file, the `server.policy` file and the `app.policy` file, then update the `app.policy` file. The symptom of a missing permission is the `java.security.AccessControlException` exception.

Note: Updates to the `app.policy` file only apply to the enterprise applications on the node to which the `app.policy` file belongs.

Windows

The missing permission is listed in the exception data, for example,

```
java.security.AccessControlException: access denied
(java.io.FilePermission
C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar read)
```

When a Java program receives this exception and adding this permission is justified, add a permission to the `server.policy` file, for example:

```
grant codeBase "file:user_client_installed_location" {
    permission java.io.FilePermission
"C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar", "read";
};
```

The previous permission information lines are split for the illustration. You actually enter the permission on one line.

To decide whether to add a permission, refer to the `AccessControlException` topic.

Restart all WebSphere Application Server enterprise applications to ensure that the updated `app.policy` file takes effect.

filter.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program. Java 2 security policy filtering is only in effect when Java 2 security is enabled.

Before modifying the `filter.policy` file, you must start the `wsadmin` tool. See the `Starting the wsadmin scripting client` article for more information.

Refer to “Protecting system resources and APIs (Java 2 security)” on page 741. The filtering policy defined in the `filter.policy` file is cell wide. The `filter.policy` file is the only policy file that is used when restricting the permission instead of granting permission. The permissions that are listed in the filter policy file are filtered out from the `app.policy` file and the `was.policy` file. Permissions that are defined in the other policy files are not affected by the `filter.policy` file.

When a permission is filtered out, an audit message is logged. However, if the permissions that are defined in the `app.policy` file and the `was.policy` file are compound permissions like the `java.security.AllPermission` permission, for example, the permission is not removed. A warning message is logged. If the Issue Permission Warning flag is enabled (default) and if the `app.policy` file and the `was.policy` file contain custom permissions (non-Java API permission, the permission package name begins with characters other than `java` or `javax`), a warning message is logged and the permission is not removed. You can change the value of the **Warn if applications are granted custom permissions** option on the Global security panel. It is not recommended that you use the `AllPermission` permission for the enterprise application.

Some default permissions that are defined in the `filter.policy` file. These permissions are the minimal ones that are recommended by the product. If more permissions are added to the `filter.policy` file, certain operations can fail for enterprise applications. Add permissions to the `filter.policy` file carefully.

You cannot use the Policy Tool to edit the `filter.policy` file. Editing must be completed in a text editor. Be careful and verify that no syntax errors exist in the `filter.policy` file. If any syntax errors exist in the `filter.policy` file, the file is not loaded by the product security runtime, which implies that filtering is disabled.

To extract the `filter.policy` file, enter the following command using information from your environment:

```
set obj [$AdminConfig extract cells/cell_name/filter.policy c:/temp/test/filter.policy]
```

To check in the policy file, enter the following command using information from your environment:

```
$AdminConfig checkin cells/cell_name/filter.policy c:/temp/test/filter.policy $obj
```

An updated `filter.policy` file is applied to all of the WebSphere Application Server enterprise applications after the servers are restarted. The `filter.policy` file is managed by configuration and file replication services.

Changes made in the file are replicated to other nodes in the cell.

The `filter.policy` file that is supplied by WebSphere Application Server resides at: `app_server_root/profiles/profile_name/config/cells/cell_name/filter.policy`.

This fill contains these permissions as defaults:

```
filterMask {
permission java.lang.RuntimePermission "exitVM";
permission java.lang.RuntimePermission "setSecurityManager";
permission java.security.SecurityPermission "setPolicy";
permission javax.security.auth.AuthPermission "setLoginConfiguration"; };
runtimeFilterMask {
permission java.lang.RuntimePermission "exitVM";
permission java.lang.RuntimePermission "setSecurityManager";
permission java.security.SecurityPermission "setPolicy";
permission javax.security.auth.AuthPermission "setLoginConfiguration"; };
```

The permissions that are defined in `filterMask` filter are for static policy filtering. The security runtime tries to remove the permissions from applications during application startup. Compound permissions are not removed, but are issued with a warning, and application deployment is stopped if applications contain permissions that are defined in the `filterMask` filter, and if scripting is used. The `runtimeFilterMask` filter defines permissions that are used by the security runtime to deny access to those permissions to application thread. Do not add more permissions to the `runtimeFilterMask` filter. Application start failure or incorrect functioning might result. Be careful when adding more permissions to the `runtimeFilterMask` filter. Usually, you only need to add permissions to the `filterMask` stanza.

WebSphere Application Server relies on the filter policy file to restrict or disallow certain permissions that can compromise the integrity of the system. For instance, WebSphere Application Server considers the `exitVM` and `setSecurityManager` permissions as those permissions that most applications never have. If these permissions are granted, the following scenarios are possible:

exitVM

A servlet, JavaServer Pages (JSP) file, enterprise bean, or other library that is used by the aforementioned might call the `System.exit` API and cause the entire WebSphere Application Server process to terminate.

setSecurityManager

An application might install its own security manager and either grant more permissions or bypass the default policy that the WebSphere Application Server security manager enforces.

Note: In application code, do not use the `setSecurityManager` permission to set a security manager. When an application uses the `setSecurityManager` permission, a conflict exists with the internal security manager within WebSphere Application Server. If you must set a security manager in an application for Remote Method Invocation (RMI) purposes, you also must select the **Use Java 2 security to restrict application access to local resources** option on the Global security panel within the WebSphere Application Server administrative console. WebSphere Application Server then registers a security manager, which the application code can verify is registered by using the `System.getSecurityManager` application programming interface (API).

For the updated `filter.policy` file to take effect, restart related Java processes.

Configuring the was.policy file:

You should update the `was.policy` file if the application has specific resources to access.

Before you begin

Java 2 security uses several policy files to determine the granted permission for each Java program. The `was.policy` file is an application-specific policy file for WebSphere Application Server enterprise applications. This file is embedded in the `META-INF/was.policy` enterprise archive (.EAR) file. The `was.policy` file is located in:

```
profile_root/config/cells/cell_name/applications/  
ear_file_name/deployments/application_name/META-INF/was.policy
```

See “Java 2 security policy files” on page 48 for the list of available policy files that are supported by WebSphere Application Server Version 6.1.

The union of the permissions that are contained in the following files is applied to the WebSphere Application Server enterprise application:

- Any policy file that is specified in the `policy.url.*` properties in the `java.security` file.
- The `app.policy` files, which are managed by configuration and file replication services.
- The `server.policy` file.
- The `java.policy` file.
- The application `was.policy` file.
- The permission specification of the `ra.xml` file.
- The shared library, which is the `library.policy` file.

Changes made in these files are replicated to other nodes in the cell.

Several product-reserved symbols are defined to associate the permission lists to a specific type of resources.

Symbol	Definition
<code>file:\${application}</code>	Permissions apply to all resources used within the application.
<code>file:\${jars}</code>	Permissions apply to all utility Java archive (JAR) files within the application
<code>file:\${ejbComponent}</code>	Permissions apply to enterprise bean resources within the application
<code>file:\${webComponent}</code>	Permissions apply to Web resources within the application
<code>file:\${connectorComponent}</code>	Permissions apply to connector resources within the application

In WebSphere Application Server, applications that manipulate threads must have the appropriate thread permissions specified in the `was.policy` or `app.policy` file. Without the thread permissions specified, the application cannot manipulate threads and WebSphere Application Server creates a `java.security.AccessControlException` exception. If you add the permissions to the `was.policy` file for a specific application, you do not need to restart WebSphere Application Server. An administrator must add the following code to a `was.policy` or `app.policy` file for an application to manipulate threads:

```
grant codeBase "file:${application}" {  
    permission java.lang.RuntimePermission "stopThread";  
    permission java.lang.RuntimePermission "modifyThread";  
    permission java.lang.RuntimePermission "modifyThreadGroup";  
};
```

An administrator can add the thread permissions to the `app.policy` file, but the permission change requires a restart of WebSphere Application Server.

Note: The Signed By and the Java Authentication and Authorization Service (JAAS) principal keywords are not supported in the `was.policy` file. The Signed By keyword is supported in the `java.policy`, `server.policy`, and `client.policy` policy file. The JAAS principal keyword is supported in a JAAS policy file when it is specified by the `java.security.auth.policy` Java virtual machine (JVM) system property. You can statically set the authorization policy files in the `java.security.auth.policy` file with the `auth.policy.url.n=URL`, where *URL* is the location of the authorization policy.

Other than these blocks, you can specify the module name for granular settings. For example,

```
grant codeBase "file:DefaultWebApplication.war" {
  permission java.security.SecurityPermission "printIdentity";
};

grant codeBase "file:IncCMP11.jar" {
  permission java.io.FilePermission
    "${user.install.root}${/}bin${/}DefaultDB${/}-",
    "read,write,delete";
};
```

Five embedded symbols are provided to specify the path and name for the `java.io.FilePermission` permission. These symbols enable flexible permission specification. The absolute file path is fixed after the application is installed.

Symbol	Definition
<code>\${app.installed.path}</code>	Path where the application is installed
<code>\${was.module.path}</code>	Path where the module is installed
<code>\${current.cell.name}</code>	Current cell name
<code>\${current.node.name}</code>	Current node name
<code>\${current.server.name}</code>	Current server name

About this task

If the default permissions for the enterprise application are enough, an action is not required. The default permissions are a union of the permissions that are defined in the `java.policy` file, the `server.policy` file, and the `app.policy` file. If an application has specific resources to access, update the `was.policy` file. The first two steps assume that you are creating a new policy file.

Note: Syntax errors in the policy files cause the application server to fail. Use care when editing these policy files.

1. Create or edit a new `was.policy` file by using the PolicyTool. For more information, see “Using PolicyTool to edit policy files” on page 743.
2. Package the `was.policy` file into the enterprise archive (EAR) file.

For more information, see “Adding the `was.policy` file to applications” on page 756. The following instructions describe how to import a `was.policy` file.

- a. Import the EAR file into an assembly tool.
- b. Open the Project Navigator view.
- c. Expand the EAR file and click **META-INF**. You might find a `was.policy` file in the META-INF directory. If you want to delete the file, right-click the file name and select **Delete**.
- d. At the bottom of the Project Navigator view, click **J2EE Hierarchy**.

- e. Import the `was.policy` file by right-clicking the **Modules** directory within the deployment descriptor and by clicking **Import > Import > File system**.
 - f. Click **Next**.
 - g. Enter the path name to the `was.policy` file in the **From directory** field or click **Browse** to locate the file.
 - h. Verify that the path directory that is listed in the **Into directory** field lists the correct META-INF directory.
 - i. Click **Finish**.
 - j. To validate the EAR file, right-click the EAR file, which contains the Modules directory, and click **Run Validation**.
 - k. To save the new EAR file, right-click the EAR file, and click **Export > Export EAR file**. If you do not save the revised EAR file, the EAR file will contain the new `was.policy` file. However, if the workspace becomes corrupted, you might lose the revised EAR file.
 - l. To generate deployment code, right-click the EAR file and click **Generate Deployment Code**.
3. Update an existing installed application, if one already exists. Modify the `was.policy` file with the Policy Tool. For more information, see “Using PolicyTool to edit policy files” on page 743.

Results

The updated `was.policy` file is applied to the application after the application restarts.

Example

When a Java program receives this exception and adding this permission is justified, add the following permission to the `was.policy` file:

To determine whether to add a permission, see “Access control exception” on page 52.

What to do next

Restart all applications for the updated `app.policy` file to take effect.

spi.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program.

For the list of available policy files that are supported by WebSphere Application Server Version 6.0.x, see “Java 2 security policy files” on page 48.

Because the default permission for the Service Provider Interface (SPI) is the AllPermission permission, the only reason to update the `spi.policy` file is a restricted SPI permission. When a change in the `spi.policy` is required, complete the following steps.

Syntax errors in the policy files cause the application server to fail. Edit these policy files carefully.

Note: Do not place the `codebase` keyword or any other keyword after the `filterMask` and `runtimeFilterMask` keywords. The `Signed By` and the Java Authentication and Authorization Service (JAAS) `Principal` keywords are not supported in the `spi.policy` file. The `Signed By` keyword is supported in the `java.policy`, `server.policy`, and `client.policy` policy files. The JAAS `Principal` keyword is supported in a JAAS policy file that is specified by the `java.security.auth.policy` Java virtual machine (JVM) system property. You can statically set the authorization policy files in `java.security.auth.policy` with `auth.policy.url.n=URL`, where `URL` is the location of the authorization policy.

To extract the `filter.policy` file, enter the following command using information from your environment:

```
set obj [$AdminConfig extract profiles/profile_name/cells/cell_name/nodes/node_name/spi.policy
c:/temp/test/spi.policy]
```

Edit the file using the Policy Tool. For more information, see “Using PolicyTool to edit policy files” on page 743.

To check in the policy file, enter the following command using information from your environment:

The updated `spi.policy` is applied to the Service Provider Interface (SPI) libraries after the Java process is restarted.

```
$AdminConfig checkin profiles/profile_name/cells/cell_name/nodes/node_name/spi.policy
c:/temp/test/spi.policy $obj
```

Examples

The `spi.policy` file is the template for SPIs or third-party resources embedded in the product. Examples of SPIs are Java Message Services (JMS) (MQSeries®) and Java database connectivity (JDBC) drivers. They are specified in the `resources.xml` file. The dynamic policy grants the permissions that are defined in the `spi.policy` file to the class paths defined in the `resources.xml` file. The union of the permission that is contained in the `java.policy` file and the `spi.policy` file are applied to the SPI libraries. The `spi.policy` files are managed by configuration and file replication services.

Changes made in these files are replicated to other nodes in the cell.

You can find the `spi.policy` file that is supplied by WebSphere Application Server in the following location: `app_server_root/profiles/profile_name/config/cells/cell_name/nodes/node_name/spi.policy`. This file contains the following default permission:

```
grant {
    permission java.security.AllPermission;
};
```

Restart the related Java processes for the changes in the `spi.policy` file to become effective.

library.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program.

For the list of available policy files that are supported by WebSphere Application Server, see “Java 2 security policy files” on page 48.

The `library.policy` file is the template for shared libraries (Java library classes). Multiple enterprise applications can define and use shared libraries. Refer to *Managing shared libraries* for information on how to define and manage the shared libraries.

If the default permissions for a shared library (union of the permissions defined in the `java.policy` file, the `app.policy` file and the `library.policy` file) are enough, no action is required. The default library policy is picked up automatically. If a specific change is required to share a library in the cell, update the `library.policy` file.

Syntax errors in the policy files cause the application server to fail. Edit these policy files carefully.

Note: Do not place the codebase keyword or any other keyword after the grant keyword. The Signed By keyword and the Java Authentication and Authorization Service (JAAS) Principal keyword are not supported in the `library.policy` file. The Signed By keyword is supported in the `java.policy`, the `server.policy`, and the `client.policy` policy files. The JAAS Principal keyword is supported in a JAAS policy file when it is specified by the Java virtual machine (JVM) system property, `java.security.auth.policy`. You can statically set the authorization policy files in the `java.security.auth.policy` file with `auth.policy.url.n=URL` where `URL` is the location of the authorization policy.

To extract the policy file, use a command prompt to enter the following command using the appropriate variable values for your environment: The previous two lines were split onto two lines for illustrative purposes only.

```
wsadmin> set obj [$AdminConfig extract cells/cell_name/nodes/  
node_name/library.policy c:/temp/test/library.policy]
```

Edit the extracted `library.policy` file with the Policy Tool. For more information, see “Using PolicyTool to edit policy files” on page 743.

To check in the policy file, use a command prompt to enter the following command using the appropriate variable values for your environment: An updated `library.policy` is applied to shared libraries after the servers restart.

```
wsadmin> $AdminConfig checkin cells/cell_name/nodes/node_name/library.policy  
c:/temp/test/library.policy $obj
```

Example

The union of the permission that is contained in the `java.policy` file, the `app.policy` file, and the `library.policy` file are applied to the shared libraries. The `library.policy` file is managed by configuration and file replication services.

Changes made in the file are replicated to other nodes in the cell.

The `library.policy` file are supplied by WebSphere Application Server resides at: `app_server_root/config/cells/cell_name/nodes/node_name/` directory. The file contains an empty permission entry as a default. For example:

```
grant {  
};
```

If the shared library in a cell requires permissions that are not defined as defaults in the `java.policy` file, the `app.policy` file and the `library.policy` file, update the `library.policy` file. The missing permission causes the `java.security.AccessControlException` exception. The missing permission is listed in the exception data.

Windows For example:

```
java.security.AccessControlException: access denied (java.io.FilePermission  
app_server_root/lib/mail-impl.jar read)
```

The previous lines are split into two lines for illustrative purposes only. The `app_server_root` variable represents your installation directory.

When a Java program receives this exception and adding this permission is justified, add a permission to the `library.policy` file.

Windows For example:

```
grant { permission java.io.FilePermission "app_server_root/lib/mail-impl.jar", "read"; };
```

The previous lines are split into two lines for illustrative purposes only. The *app_server_root* variable represents your installation directory.

To decide whether to add a permission, refer to “Access control exception” on page 52.

Restart the related Java processes for the changes in the *library.policy* file to become effective.

Adding the was.policy file to applications:

An application might need a *was.policy* file if it accesses resources that require more permissions than those granted in the default *app.policy* file.

About this task

When Java 2 security is enabled for a WebSphere Application Server, all the applications that run on WebSphere Application Server undergo a security check before accessing system resources. An application might need a *was.policy* file if it accesses resources that require more permissions than those granted in the default *app.policy* file. By default, the product security reads an *app.policy* file that is located in each node and grants the permissions in the *app.policy* file to all the applications. Include any additional required permissions in the *was.policy* file. The *was.policy* file is only required if an application requires additional permissions.

The default policy file for all applications is specified in the *app.policy* file. This file is provided by the product security, is common to all applications, and you do not change this file. Add any new permissions that are required for an application in the *was.policy* file.

The *app.policy* file supplied by WebSphere Application Server resides at *app_server_root/config/cells/profile/profile_name/config/cell_name/nodes/node_name/app.policy*. The contents of the *app.policy* file are presented in the following example:

Note: In the following code sample, the two permissions that are required by JavaMail are split onto two lines for illustration only. You actually enter the permission on one line.

```
// The following permissions apply to all the components under the application.
```

```
grant codeBase "file:${application}" {
    // The following are required by JavaMail

    permission java.io.FilePermission "
        ${was.install.root}${/}lib${/}activation-impl.jar",
    "read";

    permission java.io.FilePermission "
        ${was.install.root}${/}lib${/}mail-impl.jar", "read";

};
// The following permissions apply to all utility .jar files (other
// than enterprise beans JAR files) in the application.
grant codeBase "file:${jars}" {
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

// The following permissions apply to connector resources within the application
grant codeBase "file:${connectorComponent}" {
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
```

```

};

// The following permissions apply to all the Web modules (.war files)
// within the application.
grant codeBase "file:${webComponent}" {
    permission java.io.FilePermission "${was.module.path}${/}-", "read, write";
    // where "was.module.path" is the path where the Web module is
    // installed. Refer to Dynamic policy concepts for other symbols.
    permission java.lang.RuntimePermission "loadLibrary.*";
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

// The following permissions apply to all the EJB modules within the application.
grant codeBase "file:${ejbComponent}" {
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.util.PropertyPermission "*", "read";
};

```

If additional permissions are required for an application or for one or more modules of an application, use the `was.policy` file for that application. For example, use `codeBase` of `${application}` and add required permissions to grant additional permissions to the entire application. Similarly, use `codeBase` of `${webComponent}` and `${ejbComponent}` to grant additional permissions to all the Web modules and all the enterprise bean modules in the application. You can assign additional permissions to each module (.war file or .jar file), as shown in the following example.

This example illustrates adding extra permissions for an application in the `was.policy` file:

Note: In the following code sample, the permission for the EJB module was split onto two lines for illustration only. You actually enter the permission on one line.

```

// grant additional permissions to a Web module
grant codeBase " file:aWebModule.war" {
    permission java.security.SecurityPermission "printIdentity";
};

// grant additional permission to an EJB module
grant codeBase "file:aEJBModule.jar" {
    permission java.io.FilePermission "
        ${user.install.root}${/}bin${/}DefaultDB${/}-" ."read.write,delete";
    // where, ${user.install.root} is the system property whose value is
    // located in the app_server_root directory.
};

```

To use a `was.policy` file for your application, perform the following steps:

1. Create a `was.policy` file using the policy tool. For more information on using the policy tool, see “Using PolicyTool to edit policy files” on page 743.
2. Add the required permissions in the `was.policy` file using the policy tool.
3. Place the `was.policy` file in the application enterprise archive (EAR) file under the META-INF directory. Update the application EAR file with the newly created `was.policy` file by using the `jar` command.
4. Verify that the `was.policy` file is inserted and start an assembly tool.
5. Verify that the `was.policy` file in the application is syntactically correct. In an assembly tool, right-click the enterprise application module and click **Run Validation**.

Results

An application EAR file is now ready to run when Java 2 security is enabled.

Example

This step is required for applications to run properly when Java 2 security is enabled. If the `was.policy` file is not created and it does not contain required permissions, the application might not access system resources.

The symptom of the missing permissions is the `java.security.AccessControlException` exception. The missing permission is listed in the exception data, for example,

When an application program receives this exception and adding this permission is justified, include the permission in the `was.policy` file, for example,

The previous permission information lines are split for the illustration. Enter the permission on one line.

What to do next

Install the application.

Configuring static policy files

By configuring the static policy files, the required permission will be granted for all of the Java programs.

Before you begin

Java 2 security uses several policy files to determine the granted permission for each Java program.

See the “Java 2 security policy files” on page 48 topic for the list of available policy files that are supported by WebSphere Application Server.

Two types of policy files are supported by WebSphere Application Server: dynamic policy files and static policy files. Static policy files provide the default permissions. Dynamic policy files provide application permissions.

Policy file name	Description
<code>java.policy</code>	Contains default permissions for all of the Java programs on the node. This file seldom changes.
<code>server.policy</code>	Contains default permissions for all of the WebSphere Application Server programs on the node. This file is rarely updated.
<code>client.policy</code>	Contains default permissions for all of the applets and client containers on the node.

The static policy file is not a configuration file that is managed by the repository and the file replication service. Changes to this file are local and do not get replicated to the other machine.

1. Identify the policy file to update.
 - If the permission is required only by an application, update the dynamic policy file. Refer to “Configuring Java 2 security policy files” on page 744.
 - If the permission is required only by applets and client containers, update the `client.policy` file. Refer to “client.policy file permissions” on page 762.
 - If the permission is required only by WebSphere Application Server (servers, agents, managers and application servers), update the `server.policy` file. Refer to “server.policy file permissions” on page 760.
 - If the permission is required by all of the Java programs running on the Java virtual machine (JVM), update the `java.policy` file. Refer to “java.policy file permissions” on page 759.
2. Stop and restart WebSphere Application Server.

Results

The required permission is granted for all of the Java programs that run with the restarted JVM.

Example

If Java programs on a node require permissions, the policy file needs updating. If the Java program that required the permission is not part of an enterprise application, update the static policy file. The missing permission results in the creation of the `java.security.AccessControlException` exception. The missing permission is listed in the exception data.

For example:

```
java.security.AccessControlException: access denied (java.io.FilePermission
C:/WAS_HOME/lib/mail-impl.jar read)
```

When a Java program receives this exception and adding this permission is justified, add a permission to an adequate policy file.

For example:

```
grant codeBase "file:user_client_installed_location" {
    permission java.io.FilePermission
        "C:/WAS_HOME/lib/mail-impl.jar",
        "read";
};
```

To decide whether to add a permission, refer to “Access control exception” on page 52.

java.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program.

See “Java 2 security policy files” on page 48 for the list of available policy files that are supported by WebSphere Application Server.

The `java.policy` file is a global default policy file that is shared by all of the Java programs that run in the Java virtual machine (JVM) on the node. A change to the `java.policy` file is local for the node. The default Java policy is picked up automatically. Syntax errors in the policy files cause the application server to fail. An updated `java.policy` file is applied to all the Java programs that run in all the JVMs on the local node. Restart the programs for the updates to take effect. Modifying this file is not recommended. If a specific change is required to some of the Java programs on a node and the `java.policy` file requires updating, carefully modify the `java.policy` file with the policy tool. For more information, see “Using PolicyTool to edit policy files” on page 743.

Default permissions for the java.policy file

The `java.policy` file is not a configuration file that is managed by the repository and the file replication service. Changes to this file are local and do not get replicated to the other machine. The `java.policy` file that is supplied by WebSphere Application Server is located at `install_root/java/jre/lib/security/java.policy`. This file contains these default permissions.

```
// Standard extensions get all permissions by default
grant codeBase "file:${java.home}/lib/ext/*" {
    permission java.security.AllPermission;
};
// default permissions granted to all domains
grant {
    // Allows any thread to stop itself using the java.lang.Thread.stop()
    // method that takes no argument.
    // Note that this permission is granted by default only to remain
```

```

// backwards compatible.
// It is strongly recommended that you either remove this permission
// from this policy file or further restrict it to code sources
// that you specify, because Thread.stop() is potentially unsafe.
// See "http://java.sun.com/notes" for more information.
// permission java.lang.RuntimePermission "stopThread";

// allows anyone to listen on un-privileged ports
permission java.net.SocketPermission "localhost:1024-", "listen";

// "standard" properties that can be read by anyone

permission java.util.PropertyPermission "java.version", "read";
permission java.util.PropertyPermission "java.vendor", "read";
permission java.util.PropertyPermission "java.vendor.url", "read";
permission java.util.PropertyPermission "java.class.version", "read";
permission java.util.PropertyPermission "os.name", "read";
permission java.util.PropertyPermission "os.version", "read";
permission java.util.PropertyPermission "os.arch", "read";
permission java.util.PropertyPermission "file.separator", "read";
permission java.util.PropertyPermission "path.separator", "read";
permission java.util.PropertyPermission "line.separator", "read";

permission java.util.PropertyPermission "java.specification.version", "read";
permission java.util.PropertyPermission "java.specification.vendor", "read";
permission java.util.PropertyPermission "java.specification.name", "read";

permission java.util.PropertyPermission "java.vm.specification.version", "read";
permission java.util.PropertyPermission "java.vm.specification.vendor", "read";
permission java.util.PropertyPermission "java.vm.specification.name", "read";
permission java.util.PropertyPermission "java.vm.version", "read";
permission java.util.PropertyPermission "java.vm.vendor", "read";
permission java.util.PropertyPermission "java.vm.name", "read";
};

```

If some Java programs on a node require permissions that are not defined as defaults in the `java.policy` file, consider updating the `java.policy` file. Most of the time, other policy files are updated instead of the `java.policy` file. The missing permission causes the creation of the `java.security.AccessControlException` exception. The missing permission is listed in the exception data.

For example:

```
java.security.AccessControlException: access denied (java.io.FilePermission
C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar read)
```

The previous two lines are one continuous line.

When a Java program receives this exception and adding this permission is justified, add a permission to the `java.policy` file.

For example:

```
grant codeBase "file:user_client_installed_location" {
permission java.io.FilePermission
"C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar", "read"; };

```

To decide whether to add a permission, refer to “Access control exception” on page 52.

Restart all of the Java processes for the updated `java.policy` file to take effect.

server.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program.

The `server.policy` file is a default policy file that is shared by all of the WebSphere Application Servers on a node. The `server.policy` file is not a configuration file that is managed by the repository and the file replication service. Changes to this file are local and do not replicate to the other machine.

If the default permissions for a server (the union of the permissions that is defined in the `java.policy` file and the `server.policy` file) are enough, no action is required. The default server policy is picked up automatically. If a specific change is required to some of the server programs on a node, update the `server.policy` file with the Policy Tool. Refer to the “Using PolicyTool to edit policy files” on page 743 topic to edit policy files. Changes to the `server.policy` file are local for the node. Syntax errors in the policy files cause the application server to fail. Edit these policy files carefully. An updated `server.policy` file is applied to all the server programs on the local node. Restart the servers for the updates to take effect.

If you want to add permissions to an application, use the `app.policy` file and the `was.policy` file.

Note: Updates to the `app.policy` file only apply to the enterprise applications on the node to which the `app.policy` file belongs.

When you do need to modify the `server.policy` file, locate this file at: `profile_root/properties/server.policy`. This file contains these default permissions:

```
// Allow to use sun tools
grant codeBase "file:${java.home}/../lib/tools.jar" {
    permission java.security.AllPermission;
};

// WebSphere system classes
grant codeBase "file:${was.install.root}/plugins/-" {
    permission java.security.AllPermission;
};
grant codeBase "file:${was.install.root}/lib/-" {
    permission java.security.AllPermission;
};
grant codeBase "file:${was.install.root}/classes/-" {
    permission java.security.AllPermission;
};

// Allow the WebSphere deploy tool all permissions
grant codeBase "file:${was.install.root}/deploytool/-" {
    permission java.security.AllPermission;
};

// Allow Channel Framework classes all permission
grant codeBase "file:${was.install.root}/installedChannels/-" {
    permission java.security.AllPermission;
};

// WebSphere optional runtime classes
grant codeBase "file:${was.install.root}/optionalLibraries/-" {
    permission java.security.AllPermission;
};
```

If some server programs on a node require permissions that are not defined as defaults in the `server.policy` file and the `server.policy` file, update the `server.policy` file. The missing permission creates the `java.security.AccessControlException` exception. The missing permission is listed in the exception data.

For example:

```
java.security.AccessControlException: access denied (java.io.FilePermission
C:\WebSphere\AppServer\java\jre\lib\ext\mail-impl.jar read)
```

The previous two lines are split into two lines for illustrative purposes only.

When a Java program receives this exception and adding this permission is justified, add a permission to the `server.policy` file.

For example:

```
grant codeBase "file:user_client_installed_location" {
  permission java.io.FilePermission
  "C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar", "read"; };
```

To decide whether to add a permission, refer to “Access control exception” on page 52.

Restart all of the Java processes for the updated `server.policy` file to take effect.

client.policy file permissions:

Java 2 security uses several policy files to determine the granted permission for each Java program.

For the list of available policy files that are supported by WebSphere Application Server, see “Java 2 security policy files” on page 48.

- The `client.policy` file is a default policy file that is shared by all of the WebSphere Application Server client containers and applets on a node.
- The union of the permissions that is contained in the `java.policy` file and the `client.policy` file are given to all of the client containers for WebSphere Application Server and applets running on the node.
- The `client.policy` file is not a configuration file that is managed by the repository and the file replication service. Changes to this file are local and do not replicate to the other machine.
- The `client.policy` file supplied by WebSphere Application Server is located in the `profile_root/properties/client.policy`.
- If the default permissions for a client (union of the permissions defined in the `java.policy` file and the `client.policy` file) are enough, no action is required. The default client policy is picked up automatically.
- If a specific change is required to some of the client containers and applets on a node, modify the `client.policy` file with the Policy Tool. Refer to “Using PolicyTool to edit policy files” on page 743, to edit policy files. Changes to the `client.policy` file are local for the node.

This file contains these default permissions:

```
grant codeBase "file:${was.install.root}/java/ext/*" {
  permission java.security.AllPermission;
};

// JDK classes
grant codeBase "file:${was.install.root}/java/ext/-" {
  permission java.security.AllPermission;
};
grant codeBase "file:${was.install.root}/java/tools/ibmtools.jar" {
  permission java.security.AllPermission;
};
grant codeBase "file:/QIBM/ProdData/Java400/jdk14/lib/tools.jar" {
  permission java.security.AllPermission;
};

// WebSphere system classes
grant codeBase "file:${was.install.root}/lib/-" {
  permission java.security.AllPermission;
};
grant codeBase "file:${was.install.root}/plugins/-" {
  permission java.security.AllPermission;
};
grant codeBase "file:${was.install.root}/classes/-" {
  permission java.security.AllPermission;
};
```

```

grant codeBase "file:${was.install.root}/installedConnectors/-" {
    permission java.security.AllPermission;
};
grant codeBase "file:${user.install.root}/installedConnectors/-" {
    permission java.security.AllPermission;
};

grant codeBase "file:${was.install.root}/installedChannels/-" {
    permission java.security.AllPermission;
};

// J2EE 1.4 permissions for client container WebSphere Application Server applications
// in $WAS_HOME/installedApps
grant codeBase "file:${user.install.root}/installedApps/-" {
    //Application client permissions
    permission java.awt.AWTPermission "accessClipboard";
    permission java.awt.AWTPermission "accessEventQueue";
    permission java.awt.AWTPermission "showWindowWithoutWarningBanner";
    permission java.lang.RuntimePermission "exitVM";
    permission java.lang.RuntimePermission "loadLibrary";
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.net.SocketPermission "localhost:1024-", "accept,listen";
    permission java.io.FilePermission "*", "read,write";
    permission java.util.PropertyPermission "*", "read";
};

// J2EE 1.4 permissions for client container - expanded ear file code base
grant codeBase "file:${com.ibm.websphere.client.applicationclient.archivedir}/-" {
    permission java.awt.AWTPermission "accessClipboard";
    permission java.awt.AWTPermission "accessEventQueue";
    permission java.awt.AWTPermission "showWindowWithoutWarningBanner";
    permission java.lang.RuntimePermission "exitVM";
    permission java.lang.RuntimePermission "loadLibrary";
    permission java.lang.RuntimePermission "queuePrintJob";
    permission java.net.SocketPermission "*", "connect";
    permission java.net.SocketPermission "localhost:1024-", "accept,listen";
    permission java.io.FilePermission "*", "read,write";
    permission java.util.PropertyPermission "*", "read";
};

```

All of the client containers and applets on the local node are granted the updated permissions when they start. If some client containers or applets on a node require permissions that are not defined as defaults in the `java.policy` file and the default `client.policy` file, update the `client.policy` file. The missing permission creates the `java.security.AccessControlException` exception. The missing permission is listed in the exception data, for example,

```

java.security.AccessControlException: access denied (java.io.FilePermission
C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar read)

```

The previous two lines of sample code are one continuous line, but presented as such for illustrative purposes only.

When a client program receives this exception and adding this permission is justified, add a permission to the `client.policy` file, for example, `grant codebase "file:user_client_installed_location" { permission java.io.FilePermission "C:\WebSphere\AppServer\java\jre\lib\ext\mail.jar", "read"; }`;

To decide whether to add a permission, refer to “Access control exception” on page 52.

If you update the policy file, you must restart the browser and any client applications.

Developing with programmatic security APIs for Web applications

Use this information to programmatically secure APIs for Web applications.

Before you begin

Programmatic security is used by security-aware applications when declarative security alone is not sufficient to express the security model of the application. Programmatic security consists of the following methods of the `HttpServletRequest` interface:

getRemoteUser

Returns the user name that the client used for authentication. Returns `null` if no user is authenticated.

isUserInRole

(String role name): Returns `true` if the remote user is granted the specified security role. If the remote user is not granted the specified role, or if no user is authenticated, it returns `false`.

getUserPrincipal

Returns the `java.security.Principal` object that contains the remote user name. If no user is authenticated, it returns `null`.

You can configure several options for Web authentication that determine how the Web client interacts with protected and unprotected Uniform Resource Identifiers (URI). Also, you can specify whether WebSphere Application Server challenges the Web client for basic authentication information if the certificate authentication for the HTTPS client fails. For more information, see “Selecting an authentication mechanism” on page 247.

When the `isUserInRole` method is used, declare a `security-role-ref` element in the deployment descriptor with a `role-name` subelement containing the role name that is passed to this method, or with the `@DeclareRoles` annotation. Because actual roles are created during the assembly stage of the application, you can use a logical role as the role name and provide enough hints to the assembler in the description of the `security-role-ref` element to link that role to the actual role. During assembly, the assembler creates a `role-link` subelement to link the role name to the actual role. Creation of a `security-role-ref` element is possible if an assembly tool such as Rational Application Developer (RAD) is used. You also can create the `security-role-ref` element during assembly stage using an assembly tool.

1. Add the required security methods in the servlet code.
2. Create a `security-role-ref` element with the **role-name** field. If a `security-role-ref` element is not created during development, make sure it is created during the assembly stage.

Results

A programmatically secured servlet application.

Example

This step is required to secure an application programmatically. This action is particularly useful when a Web application needs to access external resources and wants to control the access to external resources using its own authorization table (external-resource to remote-user mapping). In this case, use the `getUserPrincipal` or the `getRemoteUser` methods to get the remote user and then it can consult its own authorization table to perform authorization. The remote user information also can help retrieve the corresponding user information from an external source such as a database or from an enterprise bean. You can use the `isUserInRole` method in a similar way.

After development, a `security-role-ref` element can be created:

```

<security-role-ref>
  <description>Provide hints to assembler for linking this role
    name to an actual role here</description>
  <role-name>Mgr</role-name>
</security-role-ref>

```

During assembly, the assembler creates a role-link element:

```

<security-role-ref>
  <description>Hints provided by developer to map the role
    name to the role-link</description>
  <role-name>Mgr</role-name>
  <role-link>Manager</role-link>
</security-role-ref>

```

You can add programmatic servlet security methods inside any servlet doGet, doPost, doPut, and doDelete service methods. The following example depicts using a programmatic security API:

```

public void doGet(HttpServletRequest request,
  HttpServletResponse response) {

  ....

  // to get remote user using getUserPrincipal()
  java.security.Principal principal = request.getUserPrincipal();
  String remoteUser = principal.getName();

  // to get remote user using getRemoteUser()
  remoteUser = request.getRemoteUser();

  // to check if remote user is granted Mgr role
  boolean isMgr = request.isUserInRole("Mgr");

  // use the above information in any way as needed by
  // the application
  ....
}

```

When developing Servlet 2.5 modules, the value of the rolename argument in isCallerInRole method can be defined using Java annotations instead of declaring a security-role-ref elements in the deployment descriptor.

```

@javax.annotation.security.DeclareRoles("Mgr")
public void doGet(HttpServletRequest request,
  HttpServletResponse response) {

  ....

  // to get remote user using getUserPrincipal()
  java.security.Principal principal = request.getUserPrincipal();
  String remoteUser = principal.getName();

  // to get remote user using getRemoteUser()
  remoteUser = request.getRemoteUser();

  // to check if remote user is granted Mgr role
  boolean isMgr = request.isUserInRole("Mgr");

  // use the above information in any way as needed by
  // the application
  ....
}

```

What to do next

After developing an application, use an assembly tool to create roles and to link the actual roles to role names in the security-role-ref elements. For more information, see *Securing Web applications using an assembly tool*.

getRemoteUser and getAuthType methods

The `getRemoteUser` and `getAuthType` methods are methods of the `javax.servlet.http.HttpServletRequest` interface. If the user has been authenticated, the `getRemoteUser` method returns the login of the user that makes the request. If the user is not authenticated, the `getRemoteUser` method returns null. The `getAuthType` method returns the name of the authentication scheme that is used to protect the servlet (for example, BASIC or SSL). If the servlet is not protected, the `getAuthType` method returns null.

For both methods, the data that is returned depends upon whether security is enabled in the application server where the servlet is deployed. The following possibilities exist:

- If security is not enabled, a servlet is requested and it is configured with Web server protection. The `getRemoteUser` method returns the login and `getAuthType` method returns the authentication scheme.
- If security is enabled and a servlet is requested, both methods return null when WebSphere Application Server protection is not configured for the servlet.
- If security is enabled, a servlet is requested, and the servlet is configured with WebSphere Application Server protection, then the `getRemoteUser` method returns the login and the `getAuthType` method returns the configured authentication scheme.

Note: You can disable security at the application server level by overriding the administrative security setting. For more information, see “Securing specific application servers” on page 70.

Example: Using a programmatic security model for a Web application

The following example depicts a Web application or servlet using the programmatic security model.

This example illustrates one use and not necessarily the only use of the programmatic security model. The application can use the information that is returned by the `getUserPrincipal`, `isUserInRole`, and the `getRemoteUser` methods in any other way that is meaningful to that application. Use the declarative security model whenever possible.

File : `HelloServlet.java`

```
public class HelloServlet extends javax.servlet.http.HttpServlet {

    public void doPost(
        javax.servlet.http.HttpServletRequest request,
        javax.servlet.http.HttpServletResponse response)
        throws javax.servlet.ServletException, java.io.IOException {
    }

    public void doGet(
        javax.servlet.http.HttpServletRequest request,
        javax.servlet.http.HttpServletResponse response)
        throws javax.servlet.ServletException, java.io.IOException {

        String s = "Hello";

        // get remote user using getUserPrincipal()
        java.security.Principal principal = request.getUserPrincipal();
        String remoteUserName = "";
        if( principal != null )
            remoteUserName = principal.getName();
        // get remote user using getRemoteUser()
        String remoteUser = request.getRemoteUser();

        // check if remote user is granted Mgr role
```

```

boolean isMgr = request.isUserInRole("Mgr");

// display Hello username for managers and bob.
if ( isMgr || remoteUserName.equals("bob") )
    s = "Hello " + remoteUserName;

String message = "<html> \n" +
    "<head><title>Hello Servlet</title></head>\n" +
    "<body> /n +"
    "<h1> " +s+ </h1>/n " +
byte[] bytes = message.getBytes();

// displays "Hello" for ordinary users
// and displays "Hello username" for managers and "bob".
response.getOutputStream().write(bytes);
}
}

```

After developing the servlet, you can create a security role reference for the HelloServlet servlet as shown in the following example:

```

<security-role-ref>
    <description> </description>
    <role-name>Mgr</role-name>
</security-role-ref>

```

Web authentication settings

Use this page to specify the Web authentication settings that are associated with a Web client.

To view this administrative console page, complete the following steps:

1. Click **Security > Global security**.
2. Under Authentication, expand **Web and SIP security** and click **General settings**.

You can override the global Web authentication setting that you select on this panel by specifying a system property on the server level. To specify the system property, complete the following steps:

1. Click **Servers > Server Types > WebSphere application servers > server_name**.
2. Under Server infrastructure, click **Java and Process Management > Process definition**.
3. Under Additional properties, click **Java Virtual Machine > Custom properties > New**

You can specify the following system properties on the server level for Web authentication.

Table 23. Web authentication system property values

Property name	Value	Explanation
com.ibm.wsspi.security.web.webAuthReq	lazy	This value is equivalent to the Authenticate only when the URI is protected option.
com.ibm.wsspi.security.web.webAuthReq	persisting	This value is equivalent to the Use available authentication data when an unprotected URI is accessed option.
com.ibm.wsspi.security.web.webAuthReq	always	This value is equivalent to the Authenticate when any URI is accessed option.
com.ibm.wsspi.security.web.failOverToBasicAuth	true	This value is equivalent to the Default to basic authentication when certificate authentication for the HTTPS client fails option.

Authenticate only when the URI is protected:

The application server challenges the Web client to provide authentication data when the Web client accesses a Uniform Resource Identifier (URI) that is protected by a Java 2 Platform, Enterprise Edition (J2EE) role. The authenticated identity is available only when the Web client accesses a protected URI.

This option is the default J2EE Web authentication behavior that is also available in previous releases of WebSphere Application Server.

Note: When you select this option, the administrative console login page is missing images. You might encounter the following error in the administrative console: "CWLAA6003: Could not display the portlet, the portlet may not be started. Check the error logs".

The missing images and the error message are a side-effect of this option. The images do not display because the URIs for the images now need authentication, which requires you to log in. You can ignore this error message.

Default: Enabled

Use available authentication data when an unprotected URI is accessed:

The Web client can access validated authenticated data that it previously could not access. This option enables the Web client to call the `getRemoteUser`, `isUserInRole`, and `getUserPrincipal` methods to retrieve an authenticated identity from an unprotected URI.

When you select this option with the **Authenticate only when the URI is protected** option, the Web client can use authenticated data when the URI is protected or not protected.

Note: This option does not challenge the Web client to provide authenticated data if the Web client accesses an unprotected URI without authenticated data.

Default: Disabled

Authenticate when any URI is accessed:

The Web client must provide authentication data regardless of whether the URI is protected.

Default: Disabled

Default to basic authentication when certificate authentication for the HTTPS client fails:

When the required HTTPS client certificate authentication fails, the application server uses the basic authentication method to challenge the Web client to provide a user ID and password.

The HTTP client certification authentication that is performed by the application server security is different from the client authentication that is performed by the Web server plug-in. If you configure the Web server plug-in for mutual authentication and client authentication fails, the following situations will occur:

- The Web server produces a error and the Web request is not processed by application server security.
- The application server cannot fail over to basic authentication.

Default: Disabled

Developing with programmatic APIs for EJB applications

Use this topic to programmatically secure your Enterprise JavaBeans (EJB) applications.

About this task

Programmatic security is used by security-aware applications when declarative security alone is not sufficient to express the security model of the application. The `javax.ejb.EJBContext` application programming interface (API) provides two methods whereby the bean provider can access security information about the enterprise bean caller.

- **isCallerInRole**(String rolename): Returns true if the bean caller is granted the security role that is specified by role name. If the caller is not granted the specified role, or if the caller is not authenticated, it returns false. If the specified role is granted Everyone access, it always returns true.
- **getCallerPrincipal**: Returns the `java.security.Principal` object that contains the bean caller name. If the caller is not authenticated, it returns a principal that contains an unauthorized name.

You can enable a login module to indicate which principal class is returned by these calls.

When the `isCallerInRole` method is used, declare a `security-role-ref` element in the deployment descriptor with a `role-name` that is subelement containing the role name that is passed to this method. Because actual roles are created during the assembly stage of the application, you can use a logical role as the role name and provide enough hints to the assembler in the description of the `security-role-ref` element to link that role to an actual role. During assembly, the assembler creates a `role-link` subelement to link the `role-name` to the actual role. Creation of a `security-role-ref` element is possible if an assembly tool such as Rational Application Developer (RAD) is used. You also can create the `security-role-ref` element during the assembly stage using an assembly tool.

1. Add the required security methods in the EJB module code.
2. Create a `security-role-ref` element with a `role-name` field for all the role names used in the `isCallerInRole` method. If a `security-role-ref` element is not created during development, make sure it is created during the assembly stage.

Results

Performing the previous steps result in a programmatically secured EJB application.

Example

Hard coding security policies in applications is strongly discouraged. The Java Platform, Enterprise Edition (Java EE) security model capabilities of declaratively specifying security policies is encouraged wherever possible. Use these APIs to develop security-aware EJB applications.

Using Java EE security model capabilities to specify security policies declaratively is useful when an EJB application wants to access external resources and wants to control the access to these external resources using its own authorization table (external-resource to user mapping). In this case, use the `getCallerPrincipal` method to get the caller identity and then the application can consult its own authorization table to perform authorization. The caller identification also can help retrieve the corresponding user information from an external source, such as database or from another enterprise bean. You can use the `isCallerInRole` method in a similar way.

After development, you can create a `security-role-ref` element:

```
<security-role-ref>
<description>Provide hints to assembler for linking this role-name to
actual role here<\description>
<role-name>Mgr<\role-name>
</security-role-ref>
```

During assembly, the assembler creates a role-link element:

```
<security-role-ref>
<description>Hints provided by developer to map role-name to role-link</description>
<role-name>Mgr</role-name>
<role-link>Manager</role-link>
</security-role-ref>
```

You can add programmatic EJB component security methods for example `isCallerInRole` and `getCallerPrincipal`, inside any business methods of an enterprise bean. The following example of programmatic security APIs includes a session bean:

```
public class aSessionBean implements SessionBean {

    ....

    // SessionContext extends EJBContext. If it is entity bean use EntityContext
    javax.ejb.SessionContext context;

    // The following method will be called by the EJB container
    // automatically
    public void setSessionContext(javax.ejb.SessionContext ctx) {
        context = ctx; // save the session bean's context
    }

    ....

    private void aBusinessMethod() {
        ....

        // to get bean's caller using getCallerPrincipal()
        java.security.Principal principal = context.getCallerPrincipal();
        String callerId= principal.getName();

        // to check if bean's caller is granted Mgr role
        boolean isMgr = context.isCallerInRole("Mgr");

        // use the above information in any way as needed by the
        //application

        ....
    }

    ....
}
```

When developing EJB 3.0 modules, the value of the `rolename` argument in `isCallerInRole` method can be defined using Java annotations instead of declaring a `security-role-ref` elements in the deployment descriptor.

```
@javax.annotation.security.DeclareRoles("Mgr")
public class aSessionBean implements SessionBean {

    ....

    // SessionContext extends EJBContext. If it is entity bean use EntityContext
    javax.ejb.SessionContext context;

    // The following method will be called by the EJB container
    // automatically
    public void setSessionContext(javax.ejb.SessionContext ctx) {
        context = ctx; // save the session bean's context
    }

}
```

```

    ....

private void aBusinessMethod() {
    ....

    // to get bean's caller using getCallerPrincipal()
    java.security.Principal principal = context.getCallerPrincipal();
    String callerId= principal.getName();

    // to check if bean's caller is granted Mgr role
    boolean isMgr = context.isCallerInRole("Mgr");

    // use the above information in any way as needed by the
    //application

    ....
}

    ....
}

```

What to do next

After developing an application, use an assembly tool to create roles and to link the actual roles to role names in the security-role-ref elements. For more information, see [Securing enterprise bean applications](#).

Example: Enterprise bean application code

The following Enterprise JavaBeans (EJB) component example illustrates the use of the `isCallerInRole` and the `getCallerPrincipal` methods in an EJB module.

Using declarative security is recommended. The following example is one way of using the `isCallerInRole` and the `getCallerPrincipal` methods. The application can use this result in any way that is suitable.

A remote interface

File : Hello.java

```

package tests;
import java.rmi.RemoteException;
/**
 * Remote interface for Enterprise Bean: Hello
 */
public interface Hello extends javax.ejb.EJBObject {
    public abstract String getMessage()throws RemoteException;
    public abstract void setMessage(String s)throws RemoteException;
}

```

A home interface

File : HelloHome.java

```

package tests;
/**
 * Home interface for Enterprise Bean: Hello
 */
public interface HelloHome extends javax.ejb.EJBHome {
    /**
     * Creates a default instance of Session Bean: Hello
     */
    public tests.Hello create() throws javax.ejb.CreateException,
        java.rmi.RemoteException;
}

```

A bean implementation

File : HelloBean.java

```
package tests;
/**
 * Bean implementation class for Enterprise Bean: Hello
 */
public class HelloBean implements javax.ejb.SessionBean {
    private javax.ejb.SessionContext mySessionCtx;
    /**
     * getSessionContext
     */
    public javax.ejb.SessionContext getSessionContext() {
        return mySessionCtx;
    }
    /**
     * setSessionContext
     */
    public void setSessionContext(javax.ejb.SessionContext ctx) {
        mySessionCtx = ctx;
    }
    /**
     * ejbActivate
     */
    public void ejbActivate() {
    }
    /**
     * ejbCreate
     */
    public void ejbCreate() throws javax.ejb.CreateException {
    }
    /**
     * ejbPassivate
     */
    public void ejbPassivate() {
    }
    /**
     * ejbRemove
     */
    public void ejbRemove() {
    }

    public java.lang.String message;

    //business methods

    // all users can call getMessage()
    public String getMessage() {
        return message;
    }

    // all users can call setMessage() but only few users can set new message.
    public void setMessage(String s) {

        // get bean's caller using getCallerPrincipal()
        java.security.Principal principal = mySessionCtx.getCallerPrincipal();
        java.lang.String callerId= principal.getName();

        // check if bean's caller is granted Mgr role
    }
}
```

```

        boolean isMgr = mySessionCtx.isCallerInRole("Mgr");

        // only set supplied message if caller is "bob" or caller is granted Mgr role
        if ( isMgr || callerId.equals("bob") )
            message = s;
        else
            message = "Hello";
    }
}

```

After the development of the entity bean, create a security role reference in the deployment descriptor under the session bean, Hello:

```

<security-role-ref>
  <description>Only Managers can call setMessage() on this bean (Hello)</description>
  <role-name>Mgr</role-name>
</security-role-ref>

```

For an explanation of how to create a <security-role-ref> element, see [Securing enterprise bean applications](#). Use the information under [Map security-role-ref and role-name to role-link](#) to create the element.

Customizing Web application login

You can create a form login page and an error page to authenticate a user.

Before you begin

A Web client or a browser can authenticate a user to a Web server using one of the following mechanisms:

- **HTTP basic authentication:** A Web server requests the Web client to authenticate and the Web client passes a user ID and a password in the HTTP header.
- **HTTPS client authentication:** This mechanism requires a user (Web client) to possess a public key certificate. The Web client sends the certificate to a Web server that requests the client certificates. This authentication mechanism is strong and uses the Hypertext Transfer Protocol with Secure Sockets Layer (HTTPS) protocol.
- **Form-based Authentication:** A developer controls the look and feel of the login screens using this authentication mechanism.

The Hypertext Transfer Protocol (HTTP) basic authentication transmits a user password from the Web client to the Web server in simple base64 encoding. Form-based authentication transmits a user password from the browser to the Web server in plain text. Therefore, both HTTP basic authentication and form-based authentication are not very secure unless the HTTPS protocol is used.

The Web application deployment descriptor contains information about which authentication mechanism to use. When form-based authentication is used, the deployment descriptor also contains entries for login and error pages. A login page can be either an HTML page or a JavaServer Pages (JSP) file. This login page displays on the Web client side when a secured resource (servlet, JSP file, HTML page) is accessed from the application. On authentication failure, an error page displays. You can write login and error pages to suit the application needs and control the look and feel of these pages. During assembly of the application, an assembler can set the authentication mechanism for the application and set the login and error pages in the deployment descriptor.

Form login uses the servlet `sendRedirect` method, which has several implications for the user. The `sendRedirect` method is used twice during form login:

- The `sendRedirect` method initially displays the form login page in the Web browser. It later redirects the Web browser back to the originally requested protected page. The `sendRedirect(String URL)` method tells the Web browser to use the HTTP GET request to get the page that is specified in the Web address. If HTTP POST is the first request to a protected servlet or JavaServer Pages (JSP) file, and no previous authentication or login occurred, then HTTP POST is not delivered to the requested page. However, HTTP GET is delivered because form login uses the `sendRedirect` method, which behaves as an HTTP GET request that tries to display a requested page after a login occurs.
 - Using HTTP POST, you might experience a scenario where an unprotected HTML form collects data from users and then posts this data to protected servlets or JSP files for processing, but the users are not logged in for the resource. To avoid this scenario, structure your Web application or permissions so that users are forced to use a form login page before the application performs any HTTP POST actions to protected servlets or JSP files.
1. Create a form login page with the required look and feel, including the required elements to perform form-based authentication. For an example, see “Example: Form login” on page 776.
 2. Create an error page. You can program error pages to retry authentication or to display an appropriate error message.
 3. Place the login page and error page in the Web archive (.war) file relative to the top directory. For example, if the login page is configured as `/login.html` in the deployment descriptor, place it in the top directory of the WAR file. An assembler can also perform this step using the assembly tool.
 4. Create a form logout page and insert it to the application only when the Web application requires a form-based authentication mechanism.

Example: Form login

See the “Example: Form login” on page 776 article for sample form login pages.

The WebSphere Application Server Samples Gallery provides a form login Sample that demonstrates how to use the WebSphere Application Server login facilities to implement and configure form login procedures. The Sample integrates the following technologies to demonstrate the WebSphere Application Server and Java Platform, Enterprise Edition (Java EE) login functionality:

- Java EE form-based login
- Java EE servlet filter with login
- IBM extension: form-based login

The form login sample is part of the Technology Samples package. For more information on how to access the form login sample, see [Accessing the Samples \(Samples Gallery\)](#).

For the authentication to proceed appropriately, the action of the login form must always have the `j_security_check` action. The following example shows how to code the form into the HTML page:

```
<form method="POST" action="j_security_check">
<input type="text" name="j_username">
<input type="text" name="j_password">
</form>
```

Use the **j_username** input field to get the user name, and use the **j_password** input field to get the user password.

On receiving a request from a Web client, the Web server sends the configured form page to the client and preserves the original request. When the Web server receives the completed form page from the Web client, the server extracts the user name and password from the form and authenticates the user. On successful authentication, the Web server redirects the call to the original request. If authentication fails, the Web server redirects the call to the configured error page.

The following example depicts a login page in HTML (`login.html`):


```

<!DOCTYPE HTML PUBLIC "-//W3C/DTD HTML 4.0 Transitional//EN">
<html>
<META HTTP-EQUIV = "Pragma" CONTENT="no-cache">
<title> Security FVT Login Page </title>
<body>
<h2>Form Login</h2>
<FORM METHOD=POST ACTION="j_security_check">
<p>
<font size="2"> <strong> Enter user ID and password: </strong></font>
<BR>
<strong> User ID</strong> <input type="text" size="20" name="j_username">
<strong> Password </strong> <input type="password" size="20" name="j_password">
<BR>
<BR>
<font size="2"> <strong> And then click this button: </strong></font>
<input type="submit" name="login" value="Login">
</p>

</form>
</body>
</html>

```

The following example depicts an error page in a JSP file:

```

<!DOCTYPE HTML PUBLIC "-//W3C/DTD HTML 4.0 Transitional//EN">
<html>
<head><title>A Form login authentication failure occurred</head></title>
<body>
<H1><B>A Form login authentication failure occurred</H1></B>
<P>Authentication may fail for one of many reasons. Some possibilities include:
<OL>
<LI>The user-id or password may be entered incorrectly; either misspelled or the
wrong case was used.
<LI>The user-id or password does not exist, has expired, or has been disabled.
</OL>
</P>
</body>
</html>

```

After an assembler configures the Web application to use form-based authentication, the deployment descriptor contains the login configuration as shown:

```

<login-config id="LoginConfig_1">
<auth-method>FORM</auth-method>
<realm-name>Example Form-Based Authentication Area</realm-name>
<form-login-config id="FormLoginConfig_1">
<form-login-page>/login.html</form-login-page>
<form-error-page>/error.jsp</form-error-page>
</form-login-config>
</login-config>

```

A sample Web application archive (WAR) file directory structure that shows login and error pages for the previous login configuration follows:

```

META-INF
  META-INF/MANIFEST.MF
  login.html
  error.jsp
WEB-INF/
  WEB-INF/classes/
  WEB-INF/classes/aServlet.class

```

What to do next

After developing login and error pages, add them to the Web application. Use the assembly tool to configure an authentication mechanism and insert the developed login page and error page in the deployment descriptor of the application.

Example: Form login

This article provides several examples pertaining to form login.

For the authentication to proceed appropriately, the action of the login form must always have the `j_security_check` action. The following example shows how to code the form into the HTML page:

```
<form method="POST" action="j_security_check">
<input type="text" name="j_username">
<input type="text" name="j_password">
</form>
```

Use the `j_username` input field to get the user name, and use the `j_password` input field to get the user password.

On receiving a request from a Web client, the Web server sends the configured form page to the client and preserves the original request. When the Web server receives the completed form page from the Web client, the server extracts the user name and password from the form and authenticates the user. On successful authentication, the Web server redirects the call to the original request. If authentication fails, the Web server redirects the call to the configured error page.

The following example depicts a login page in HTML (`login.html`):

```
<!DOCTYPE HTML PUBLIC "-//W3C/DTD HTML 4.0 Transitional//EN">
<html>
<META HTTP-EQUIV = "Pragma" CONTENT="no-cache">
<title> Security FVT Login Page </title>
<body>
<h2>Form Login</h2>
<FORM METHOD=POST ACTION="j_security_check">
<p>
<font size="2"> <strong> Enter user ID and password: </strong></font>
<BR>
<strong> User ID</strong> <input type="text" size="20" name="j_username">
<strong> Password </strong> <input type="password" size="20" name="j_password">
<BR>
<BR>
<font size="2"> <strong> And then click this button: </strong></font>
<input type="submit" name="login" value="Login">
</p>

</form>
</body>
</html>
```

The following example depicts an error page in a JSP file:

```
<!DOCTYPE HTML PUBLIC "-//W3C/DTD HTML 4.0 Transitional//EN">
<html>
<head><title>A Form login authentication failure occurred</head></title>
<body>
<H1><B>A Form login authentication failure occurred</H1></B>
<P>Authentication may fail for one of many reasons. Some possibilities include:
<OL>
<LI>The user-id or password may be entered incorrectly; either misspelled or the
wrong case was used.
<LI>The user-id or password does not exist, has expired, or has been disabled.
```

```
</OL>
</P>
</body>
</html>
```

After an assembler configures the Web application to use form-based authentication, the deployment descriptor contains the login configuration as shown:

```
<login-config id="LoginConfig_1">
<auth-method>FORM</auth-method>
<realm-name>Example Form-Based Authentication Area</realm-name>
<form-login-config id="FormLoginConfig_1">
<form-login-page>/login.html</form-login-page>
<form-error-page>/error.jsp</form-error-page>
</form-login-config>
</login-config>
```

A sample Web application archive (WAR) file directory structure that shows login and error pages for the previous login configuration follows:

```
META-INF
  META-INF/MANIFEST.MF
  login.html
  error.jsp
WEB-INF/
  WEB-INF/classes/
  WEB-INF/classes/aServlet.class
```

Form logout

Form logout is a mechanism to log out without having to close all Web-browser sessions. After logging out of the form logout mechanism, access to a protected Web resource requires re-authentication. This feature is not required by J2EE specifications, but it is provided as an additional feature in WebSphere Application Server security.

Suppose that you want to log out after logging into a Web application and perform some actions. A form logout works in the following manner:

1. The logout-form URI is specified in the Web browser and loads the form.
2. The user clicks **Submit** on the form to log out.
3. The WebSphere Application Server security code logs the user out. During this process, the Application Server completes the following processes:
 - a. Clears the Lightweight Third Party Authentication (LTPA) / single sign-on (SSO) cookies
 - b. Invalidates the HTTP session
 - c. Removes the user from the authentication cache
4. Upon logout, the user is redirected to a logout exit page.

Form logout does not require any attributes in a deployment descriptor. The form-logout page is an HTML or a JavaServer Pages (JSP) file that is included with the Web application. The form-logout page is like most HTML forms except that like the form-login page, the form-logout page has a special post action. This post action is recognized by the Web container, which dispatches the post action to a special internal form-logout servlet. The post action in the form-logout page must be `ibm_security_logout`.

You can specify a logout-exit page in the logout form and the exit page can represent an HTML or a JSP file within the same Web application to which the user is redirected after logging out. Additionally, the logout-exit page permits a fully qualified URL in the form of `http://hostname:port/URL`. The logout-exit page is specified as a parameter in the form-logout page. If no logout-exit page is specified, a default logout HTML message is returned to the user.

Here is a sample form logout HTML form. This form configures the logout-exit page to redirect the user back to the login page after logout.

```

<!DOCTYPE HTML PUBLIC "-//W3C/DTD HTML 4.0 Transitional//EN">
<html>
  <META HTTP-EQUIV = "Pragma" CONTENT="no-cache">
  <title>Logout Page </title>
  <body>
    <h2>Sample Form Logout</h2>
    <FORM METHOD=POST ACTION="ibm_security_logout" NAME="logout">
      <p>
        <BR>
        <BR>
        <font size="2"><strong> Click this button to log out: </strong></font>
        <input type="submit" name="logout" value="Logout">
        <INPUT TYPE="HIDDEN" name="logoutExitPage" VALUE="/login.html">
      </p>
    </form>
  </body>
</html>

```

The WebSphere Application Server Samples Gallery provides a form login Sample that demonstrates how to use the WebSphere Application Server login facilities to implement and configure form login procedures. The Sample integrates the following technologies to demonstrate the WebSphere Application Server and Java 2 Platform, Enterprise Edition (J2EE) login functionality:

- J2EE form-based login
- J2EE servlet filter with login
- IBM extension: form-based login

The form login Sample is part of the Technology Samples package.

Developing servlet filters for form login processing

You can control the look and feel of the login screen using the form-based login mechanism. In form-based login, you specify a login page that is used to retrieve the user ID and password information. You also can specify an error page that displays when authentication fails.

About this task

If additional authentication or additional processing is required before and after authentication, servlet filters are an option. Servlet filters can dynamically intercept requests and responses to transform or to use the information that is contained in the requests or responses. One or more servlet filters can be attached to a servlet or to a group of servlets. Servlet filters also can attach to JavaServer Pages (JSP) files and HTML pages. All of the attached servlet filters are called before the servlet is invoked.

Both form-based login and servlet filters are supported by any servlet Version 2.3 specification-complaint Web container. The form login servlet performs the authentication and servlet filters perform additional authentication, auditing, or logging information.

To perform pre-login and post-login actions using servlet filters, configure these filters for either form login page support or for the `/j_security_check` URL. The `j_security_check` is posted by a form login page with the `j_username` parameter that contains the user name and the `j_password` parameter that contains the password. A servlet filter can use the user name parameter and password information to perform more authentication or other special needs.

1. A servlet filter implements the `javax.servlet.Filter` class. Implement three methods in the filter class:
 - **init(javax.servlet.FilterConfig cfg)**. This method is called by the container once, when the servlet filter is placed into service. The `FilterConfig` passed to this method contains the init-parameters of the servlet filter. Specify the init-parameters for a servlet filter during configuration using the assembly tool.
 - **destroy**. This method is called by the container when the servlet filter is taken out of a service.

- **doFilter(ServletRequest req, ServletResponse res, FilterChain chain).** This method is called by the container for every servlet request that maps to this filter before invoking the servlet. The FilterChain chain that is passed to this method can be used to invoke the next filter in the chain of filters. The original requested servlet runs when the last filter in the chain calls the chain.doFilter method. Therefore, all filters call the chain.doFilter method for the original servlet to run after filtering. If an additional authentication check is implemented in the filter code and results in failure, the original servlet does not run. The chain.doFilter method is not called and can be redirected to some other error page.
2. If a servlet maps to many servlet filters, servlet filters are called in the order that is listed in the web.xml deployment descriptor of the application. Place the servlet filter class file in the WEB-INF/classes directory of the application.

Example

An example of a servlet filter follows: This login filter can map to the /j_security_check URL to perform pre-login and post-login actions.

```
import javax.servlet.*;
public class LoginFilter implements Filter {
    protected FilterConfig filterConfig;
    // Called once when this filter is instantiated.
    // If mapped to j_security_check, called
    // very first time j_security_check is invoked.
    public void init(FilterConfig filterConfig) throws ServletException {
        this.filterConfig = filterConfig;
    }
    public void destroy() {
        this.filterConfig = null;
    }
    // Called for every request that is mapped to this filter.
    // If mapped to j_security_check,
    // called for every j_security_check action
    public void doFilter(ServletRequest request,
        ServletResponse response, FilterChain chain)
        throws java.io.IOException, ServletException {
        // perform pre-login action here
        chain.doFilter(request, response);
        // calls the next filter in chain.
        // j_security_check if this filter is
        // mapped to j_security_check.
        // perform post-login action here.
    }
}
```

Example: Using servlet filters to perform pre-login and post-login processing during form login

This example illustrates one way that the servlet filters can perform pre-login and post-login processing during form login.

Servlet filter source code: LoginFilter.java

```
/**
 * A servlet filter example: This example filters j_security_check and
 * performs pre-login action to determine if the user trying to log in
 * is in the revoked list. If the user is on the revoked list, an error is
 * sent back to the browser.
 *
 * This filter reads the revoked list file name from the FilterConfig
 * passed in the init() method. It reads the revoked user list file and
 * creates a revokedUsers list.
 *
 * When the doFilter method is called, the user logging in is checked
 * to make sure that the user is not on the revoked Users list.
 */
import javax.servlet.*;
import javax.servlet.http.*;
```

```

import java.io.*;

public class LoginFilter implements Filter {

    protected FilterConfig filterConfig;

    java.util.List revokeList;

    /**
     * init() : init() method called when the filter is instantiated.
     * This filter is instantiated the first time j_security_check is
     * invoked for the application (When a protected servlet in the
     * application is accessed).
     */
    public void init(FilterConfig filterConfig) throws ServletException {
        this.filterConfig = filterConfig;

        // read revoked user list
        revokeList = new java.util.ArrayList();
        readConfig();
    }

    /**
     * destroy() : destroy() method called when the filter is taken
     * out of service.
     */
    public void destroy() {
        this.filterConfig = null;
        revokeList = null;
    }

    /**
     * doFilter() : doFilter() method called before the servlet to
     * which this filter is mapped is invoked. Since this filter is
     * mapped to j_security_check, this method is called before
     * j_security_check action is posted.
     */
    public void doFilter(ServletRequest request, ServletResponse response,
        FilterChain chain) throws java.io.IOException, ServletException {

        HttpServletRequest req = (HttpServletRequest)request;
        HttpServletResponse res = (HttpServletResponse)response;

        // pre login action

        // get username
        String username = req.getParameter("j_username");

        // if user is in revoked list send error
        if ( revokeList.contains(username) ) {
            res.sendError(javax.servlet.http.HttpServletResponse.SC_UNAUTHORIZED);
            return;
        }

        // call next filter in the chain : let j_security_check authenticate
        // user
        chain.doFilter(request, response);

        // post login action
    }

    /**
     * readConfig() : Reads revoked user list file and creates a revoked
     * user list.
     */
    private void readConfig() {
        if ( filterConfig != null ) {

```

```

// get the revoked user list file and open it.
BufferedReader in;
try {
    String filename = filterConfig.getInitParameter("RevokedUsers");
    in = new BufferedReader( new FileReader(filename));
} catch ( FileNotFoundException fnfe) {
    return;
}

// read all the revoked users and add to revokeList.
String userName;
try {
    while ( (userName = in.readLine()) != null )
        revokeList.add(userName);
} catch ( IOException ioe) {
}

}

}

```

Note: In the previous code sample, the line that begins `public void doFilter(ServletRequest request` is broken into two lines for illustrative purposes only. The `public void doFilter(ServletRequest request` line and the line after it are one continuous line.

An example of the `web.xml` file that shows the `LoginFilter` filter configured and mapped to the `j_security_check` URL:

```

<filter id="Filter_1">
  <filter-name>LoginFilter</filter-name>
  <filter-class>LoginFilter</filter-class>
  <description>Performs pre-login and post-login operation</description>
  <init-param>
    <param-name>RevokedUsers</param-name>
    <param-value>c:\WebSphere\AppServer\installedApps\
      <app-name>\revokedUsers.lst</param-value>
  </init-param>
</filter-id>

<filter-mapping>
  <filter-name>LoginFilter</filter-name>
  <url-pattern>/j_security_check</url-pattern>
</filter-mapping>

```

An example of a revoked user list file:

```

user1
cn=user1,o=ibm,c=us
user99
cn=user99,o=ibm,c=us

```

Configuring servlet filters

IBM Rational Application Developer or an assembly tool can configure the servlet filters. Two steps are involved in configuring a servlet filter.

1. Name the servlet filter and assign the corresponding implementation class to the servlet filter.

Optionally, assign initialization parameters that get passed to the `init` method of the servlet filter. After configuring the servlet filter, the `web.xml` application deployment descriptor contains a servlet filter configuration similar to the following example:

```

<filter id="Filter_1">
  <filter-name>LoginFilter</filter-name>
  <filter-class>LoginFilter</filter-class>

```



```

<description>Performs pre-login and post-login
    operation</description>
<init-param>// optional
    <param-name>ParameterName</param-name>
    <param-value>ParameterValue</param-value>
</init-param>
</filter>

```

2. Map the servlet filter to a URL or a servlet.

After mapping the servlet filter to a URL or a servlet, the `web.xml` application deployment descriptor contains servlet mapping similar to the following example:

```

<filter-mapping>
    <filter-name>LoginFilter</filter-name>
    <url-pattern>/j_security_check</url-pattern>
    // can be servlet <servlet>servletName</servlet>
</filter-mapping>

```

Example

You can use servlet filters to replace the `CustomLoginServlet` servlet, and to perform additional authentication, auditing, and logging.

The WebSphere Application Server Samples Gallery provides a form login sample that demonstrates how to use the WebSphere Application Server login facilities to implement and configure form login procedures. The sample integrates the following technologies to demonstrate the WebSphere Application Server and Java Platform, Enterprise Edition (Java EE) login functionality:

- Java EE form-based login
- Java EE servlet filter with login
- IBM extension: form-based login

The form login sample is part of the Technology Samples package. For more information on how to access the form login sample, see [Accessing the Samples \(Samples Gallery\)](#).

Secure transports with JSSE and JCE programming interfaces

This topic provides detailed information about transport security using Java Secure Socket Extension (JSSE) and Java Cryptography Extension (JCE) programming interfaces. Within this topic, there is a description of the IBM version of the Java Cryptography Extension Federal Information Processing Standard (IBMJCE/FIPS).

Java Secure Socket Extension

Java Secure Socket Extension (JSSE) provides the transport security for WebSphere Application Server. JSSE provides the application programming interface (API) framework and the implementation of the APIs for Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols, including functionality for data encryption, message integrity, and authentication.

JSSE APIs are integrated into the Java 2 SDK, Standard Edition (J2SDK), Version 5. The API package for JSSE APIs is `javax.net.ssl.*`. Documentation for using JSSE APIs can be found in the J2SE 6 API documentation that is located at <http://java.sun.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html>.

Several JSSE providers ship with the Java 2 SDK Version 5 that comes with WebSphere Application Server. The IBMJSSE provider is used in previous WebSphere Application Server releases. Associated with the IBMJSSE provider is the IBMJSSE/FIPS provider, which is used when FIPS is enabled on the

server. Both of these providers do not work with the Java Message Service (JMS) and HTTP transports in WebSphere Application Server Version 7.0. These transports take advantage of the J2SDK Verison 5 network input/output (NIO) asynchronous channels.

For more information on the new IBMJSSE2 provider, please review the documentation located at <http://www.ibm.com/developerworks/java/jdk/security/60/>.

Customizing Java Secure Socket Extension

You can customize a number of aspects of JSSE by plugging in different implementations of Cryptography Package Provider, X509Certificate and HTTPS protocols, or specifying different default keystore files, key manager factories, and trust manager factories. The following table summarizes which aspects can be customized, what the defaults are, and which mechanisms are used to provide customization. You can customize the following key aspects:

Customizable item	Default	How to customize
X509Certificate	X509Certificate implementation from IBM	The <code>cert.provider.x509v1</code> security property
HTTPS protocol	Implementation from IBM	The <code>java.protocol.handler.pkgs</code> system property
Cryptography Package Provider	IBMJSSE2	A <code>security.provider.n=</code> line in security properties file. See description.
Default keystore	None	The <code>* javax.net.ssl.keyStore</code> system property
Default truststore	<code>jssecacerts</code> , if it exists. Otherwise, <code>cacerts</code>	The <code>* javax.net.ssl.trustStore</code> system property
Default key manager factory	<code>IbmX509</code>	The <code>ssl.KeyManagerFactory.algorithm</code> security property
Default trust manager factory	<code>IbmX509</code>	The <code>ssl.TrustManagerFactory.algorithm</code> security property

For aspects that you can customize by setting a system property, statically set the system property by using the **-D** option of the **Java** command. You can set the system property using the administrative console, or set the system property dynamically by calling the `java.lang.System.setProperty` method in your code: `System.setProperty(propertyName, "propertyValue")`.

For aspects that you can customize by setting a Java security property, statically specify a security property value in the `java.security` properties file. The security property is `propertyName=propertyValue`. Dynamically set the Java security property by calling the `java.security.Security.setProperty` method in your code.

The `java.security` properties file is located in the following directory:

`app_server_root/java/jre/lib/security` directory.

Application Programming Interface

The JSSE provides a standard application programming interface (API) that is available in packages of the `javax.net` file, `javax.net.ssl` file, and the `javax.security.cert` file. The APIs cover:

- Sockets and SSL sockets
- Factories to create the sockets and SSL sockets
- Secure socket context that acts as a factory for secure socket factories
- Key and trust manager interfaces
- Secure HTTP URL connection classes
- Public key certificate API

You can find more information documented for the JSSE APIs if you access the following information:

Version 1.6

1. Access the <http://www.ibm.com/developerworks/java/jdk/security/> Web site.
2. Click **Java 1.6**.
3. Click **Javadoc HTML documentation** in the Java Secure Socket Extension (JSSE) Guide section.

Samples using Java Secure Socket Extension

The Java Secure Socket Extension (JSSE) also provides samples to demonstrate its functionality. The Java Secure Socket Extension (JSSE) also provides samples to demonstrate its functionality. You can access the samples in the following location:

Version 1.6

1. Access the <http://www.ibm.com/developerworks/java/jdk/security/> Web site.
2. Click **Java 1.6**.
3. Click **jssedocs_samples.zip** in the Java Secure Socket Extension (JSSE) Guide section.

Look for the following files:

Files	Description
ClientJsse.java	Demonstrates a simple client and server interaction using JSSE. All enabled cipher suites are used.
OldServerJsse.java	Back-level samples
ServerPKCS12Jsse.java	Demonstrates a simple client and server interaction using JSSE with the PKCS12 keystore file. All enabled cipher suites are used.
ClientPKCS12Jsse.java	Demonstrates a simple client and server interaction using JSSE with the PKCS12 keystore file. All enabled cipher suites are used.
UseHttps.java	Demonstrates accessing an SSL or non-SSL Web server using the Java protocol handler of the <code>com.ibm.net.ssl.www.protocol</code> class. The URL is specified with the <code>http</code> or <code>https</code> prefix. The HTML that is returned from this site is displayed.

See more instructions in the source code. Follow these instructions before you run the samples.

Permissions for Java 2 security

You might need the following permissions to run an application with JSSE: This list is for reference only.

- `java.util.PropertyPermission "java.protocol.handler.pkgs", "write"`
- `java.lang.RuntimePermission "writeFileDescriptor"`
- `java.lang.RuntimePermission "readFileDescriptor"`
- `java.lang.RuntimePermission "accessClassInPackage.sun.security.x509"`
- `java.io.FilePermission "${user.install.root}/${etc}/.keystore", "read"`
- `java.io.FilePermission "${user.install.root}/${etc}/.truststore", "read"`

For the IBMJSSE provider:

- `java.security.SecurityPermission "putProviderProperty.IBMJSSE"`
- `java.security.SecurityPermission "insertProvider.IBMJSSE"`

For the SUNJSSE provider:

- `java.security.SecurityPermission "putProviderProperty.SunJSSE"`
- `java.security.SecurityPermission "insertProvider.SunJSSE"`

Debugging

By configuring through the `javax.net.debug` system property, JSSE provides the following dynamic debug tracing: `-Djavax.net.debug=true`.

A value of `true` turns on the trace facility, provided that the debug version of JSSE is installed.

Documentation

See the Security: Resources for learning topic for documentation references to JSSE.

JCE

Java Cryptography Extension (JCE) provides cryptographic, key and hash algorithms for WebSphere Application Server. JCE provides a framework and implementations for encryption, key generation, key agreement, and Message Authentication Code (MAC) algorithms. Support for encryption includes symmetric, asymmetric, block and stream ciphers.

IBMJCE

The IBM version of the Java Cryptography Extension (IBMJCE) is an implementation of the JCE cryptographic service provider that is used in WebSphere Application Server. The IBMJCE is similar to SunJCE, except that the IBMJCE offers more algorithms:

- Cipher algorithm (AES, DES, TripleDES, PBEs, Blowfish, and so on)
- Signature algorithm (SHA1withRSA, MD5withRSA, SHA1withDSA)
- Message digest algorithm (MD5, MD2, SHA1, SHA-256, SHA-384, SHA-512)
- Message authentication code (HmacSHA1, HmacMD5)
- Key agreement algorithm (DiffieHellman)
- Random number generation algorithm (IBMSecureRandom, SHA1PRNG)
- Key store (JKS, JCEKS, PKCS12, JCERACFKS [z/OS only])

The IBMJCE belongs to the `com.ibm.crypto.provider.*` packages.

For further information, see the information on JCE on the following web site: <http://www.ibm.com/developerworks/java/jdk/security/60/>.

IBMJCEFIPS

The IBM version of the Java Cryptography Extension Federal Information Processing Standard (IBMJCEFIPS) is an implementation of the JCE cryptographic service provider that is used in WebSphere Application Server. The IBMJCEFIPS service provider implements the following:

- Signature algorithms (SHA1withDSA, SHA1withRSA)
- Cipher algorithms (AES, TripleDES, RSA)
- Key agreement algorithm (DiffieHellman)
- Key (pair) generator (DSA, AES, TripleDES, HmacSHA1, RSA, DiffieHellman)
- Message authentication code (MAC) (HmacSHA1)
- Message digest (MD5, SHA-1, SHA-256, SHA-384, SHA-512)
- Algorithm parameter generator (DiffieHellman, DSA)
- Algorithm parameter (AES, DiffieHellman, DES, TripleDES, DSA)
- Key factory (DiffieHellman, DSA, RSA)
- Secret key factory (AES, TripleDES)
- Certificate (X.509)
- Secure random (IBMSecureRandom)

Application Programming Interface

Java Cryptography Extension (JCE) has a provider-based architecture. Providers can be plugged into the JCE framework by implementing the APIs that are defined by the JCE. The JCE APIs cover:

- Symmetric bulk encryption, such as DES, RC2, and IDEA
- Symmetric stream encryption, such as RC4
- Asymmetric encryption, such as RSA
- Password-based encryption (PBE)
- Key agreement
- Message authentication codes

There is more information documented for the JCE APIs on the <http://www.ibm.com/developerworks/java/jdk/security/> Web site.

Samples using Java Cryptography Extension

There are samples located on the <http://www.ibm.com/developerworks/java/jdk/security/> Web site in the `jceDocs_samples.zip` file. Unzip the file and locate the following samples in the `jceDocs/samples` directory:

File	Description
<code>SampleDSASignature.java</code>	Demonstrates how to generate a pair of DSA keys (a public key and a private key) and use the key to digitally sign a message using the SHA1withDSA algorithm
<code>SampleMarsCrypto.java</code>	Demonstrates how to generate a Mars secret key, and how to do Mars encryption and decryption
<code>SampleMessageDigests.java</code>	Demonstrates how to use the message digest for MD2 and MD5 algorithms
<code>SampleRSACrypto.java</code>	Demonstrates how to generate an RSA key pair, and how to do RSA encryption and decryption
<code>SampleRSASignatures.java</code>	Demonstrates how to generate a pair of RSA keys (a public key and a private key) and use the key to digitally sign a message using the SHA1withRSA algorithm
<code>SampleX509Verification.java</code>	Demonstrates how to verify X509 certificates

Documentation

Refer to the Security: Resources for learning for documentation on JCE.

Configuring Federal Information Processing Standard Java Secure Socket Extension files

Use this topic to configure Federal Information Processing Standard Java Secure Socket Extension files.

About this task

In WebSphere Application Server, the Java Secure Socket Extension (JSSE) provider used is the IBMJSSE2 provider. This provider delegates encryption and signature functions to the Java Cryptography Extension (JCE) provider. Consequently, IBMJSSE2 does not need to be Federal Information Processing Standard (FIPS)-approved because it does not perform cryptography. However, the JCE provider requires FIPS-approval.

WebSphere Application Server provides a FIPS-approved IBMJCEFIPS provider that IBMJSSE2 can utilize. The IBMJCEFIPS provider that is shipped in WebSphere Application Server Version 7.0 supports the following SSL ciphers:

- SSL_RSA_WITH_AES_128_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_RSA_WITH_AES_128_CBC_SHA
- SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_DSS_WITH_AES_128_CBC_SHA
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA

Even though the IBMJSSEFIPS provider is still present, the runtime does not use this provider. If IBMJSSEFIPS is specified as a contextProvider, WebSphere Application Server automatically defaults to the IBMJSSE2 provider (with the IBMJCEFIPS provider) for supporting FIPS. When enabling the **Use the United States Federal Information Processing Standard (FIPS) algorithms** option on the server SSL certificate and key management panel, the runtime always uses IBMJSSE2, despite the contextProvider that you specify for SSL (IBMJSSE, IBMJSSE2 or IBMJSSEFIPS). Also, because FIPS requires the SSL protocol be TLS, the runtime always uses TLS when FIPS is enabled, regardless of the SSL protocol setting in the SSL repertoire. This simplifies the FIPS configuration in Version 7.0 because an administrator needs to enable only the **Use the United States Federal Information Processing Standard (FIPS) algorithms** option on the server SSL certificate and key management panel to enable all transports using SSL.

1. Click **Security > SSL certificate and key management**.
2. Select the **Use the United States Federal Information Processing Standard (FIPS) algorithms** option and click **Apply**. This option makes IBMJSSE2 and IBMJCEFIPS the active providers.
3. Accommodate Java clients that must access enterprise beans.

Change the `com.ibm.security.useFIPS` property value from `false` to `true` in the `profile_root/properties/ssl.client.props` file.

4. Ensure that the `java.security` includes the provider.

Edit the `java.security` file to insert the IBMJCEFIPS provider (`com.ibm.crypto.fips.provider.IBMJCEFIPS`) before the IBMJCE provider, and also renumber the other providers in the provider list. The IBMJCEFIPS provider must be in the `java.security` file provider list.

The `java.security` file is located in the `WASHOME/java/jre/lib/security` directory.

The IBM SDK `java.security` file looks like the following example after completing this step:

```
security.provider.1=com.ibm.crypto.fips.provider.IBMJCEFIPS
security.provider.2=com.ibm.crypto.provider.IBMJCE
security.provider.3=com.ibm.jsse.IBMJSSEProvider
security.provider.4=com.ibm.jsse2.IBMJSSEProvider2
security.provider.5=com.ibm.security.jgss.IBMJGSSProvider
security.provider.6=com.ibm.security.cert.IBMCertPath
security.provider.7=com.ibm.crypto.pkcs11.provider.IBMPKCS11
security.provider.8=com.ibm.security.cmskeystore.CMSProvider
security.provider.9=com.ibm.security.jgss.mech.spnego.IBMPNEGO
```

If you are using the Sun Java SE Development Kit, the `java.security` file looks like the following example after completing this step:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.crypto.fips.provider.IBMJCEFIPS
security.provider.3=com.ibm.crypto.provider.IBMJCE
security.provider.4=com.ibm.jsse.IBMJSSEProvider
security.provider.5=com.ibm.jsse2.IBMJSSEProvider2
security.provider.6=com.ibm.security.jgss.IBMJGSSProvider
security.provider.7=com.ibm.security.cert.IBMCertPath
```



```
security.provider.8=com.ibm.i5os.jsse.JSSEProvider
#security.provider.8=com.ibm.crypto.pkcs11.provider.IBMPKCS11
security.provider.9=com.ibm.security.jgss.mech.spnego.IBMSPNEGO
```

What to do next

After completing these steps, a FIPS-approved JSSE or JCE provider offers increased encryption capabilities. However, when you use FIPS-approved providers:

- By default, Microsoft Internet Explorer might not have TLS enabled. To enable TLS, open the Internet Explorer browser and click **Tools > Internet Options**. On the Advanced tab, select the Use TLS 1.0 option.

Note: Netscape Version 4.7.x and earlier versions might not support TLS.

- IBM Directory Server Version 5.1 (and earlier versions) do not support TLS.
- If you have an administrative client that uses a SOAP connector and you enable FIPS, add the following line to the *profile_root/properties/soap.client.props* file:

```
com.ibm.ssl.contextProvider=IBMJSSEFIPS
```

- When you select the **Use the Federal Information Processing Standard (FIPS)** option on the SSL certificate and key management panel, the Lightweight Third-Party Authentication (LTPA) token format is not backwards-compatible with previous releases of WebSphere Application Server. However, you can import the LTPA keys from a previous version of the application server.

Note: The following error might occur when you attempt to stop WebSphere Application Server after enabling the FIPS option:

```
ADMU3007E: Exception com.ibm.websphere.management.exception.ConnectorException
```

Uncomment the following entry in the *java.security* file if it was previously removed or commented out, then restart the server:

```
security.provider.2=com.ibm.crypto.provider.IBMJCE
```

The *java.security* file is located in the *WAS_HOME/java/jre/lib/security* directory.

Note: When enabling FIPS, you cannot configure cryptographic token devices in the SSL repertoires. IBMJSSE2 must use IBMJCEFIPS when utilizing cryptographic services for FIPS.

The following FIPS 140-2 approved cryptographic providers that are the only devices that are supported with the FIPS option:

- IBMJCEFIPS (certificate 376)
- IBM Cryptography for C (ICC) (certificate 384)

The relevant certificates are listed on the NIST Web site: Cryptographic Module Validation Program FIPS 140-1 and FIPS 140-2 Pre-validation List

To unconfigure the FIPS provider, reverse the changes that you made in the previous steps. After you reverse the changes, verify that you have made the following changes to the *sas.client.props*, *soap.client.props*, and *java.security* files:

- In the *ssl.client.props* file, you must change the *com.ibm.security.useFIPS* value to false.
- In the *java.security* file, you must change the FIPS provider to a non-FIPS provider.

If you are using the IBM SDK *java.security* file, you must change the first provider to a non-FIPS provider as shown in the following example:

```
#security.provider.1=com.ibm.crypto.fips.provider.IBMJCEFIPS
security.provider.1=com.ibm.crypto.provider.IBMJCE
security.provider.2=com.ibm.jsse.IBMJSSEProvider
security.provider.3=com.ibm.jsse2.IBMJSSEProvider2
```



```
security.provider.4=com.ibm.security.jgss.IBMJGSSProvider
security.provider.5=com.ibm.security.cert.IBMCertPath
#security.provider.6=com.ibm.crypto.pkcs11.provider.IBMPKCS11
```

If you are using the Sun JDK `java.security` file, you must change the third provider to a non-FIPS provider as shown in the following example:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.security.jgss.IBMJGSSProvider
security.provider.3=com.ibm.crypto.fips.provider.IBMJCEFIPS
security.provider.4=com.ibm.crypto.provider.IBMJCE
security.provider.5=com.ibm.jsse.IBMJSSEProvider
security.provider.6=com.ibm.jsse2.IBMJSSEProvider2
security.provider.7=com.ibm.security.cert.IBMCertPath
#security.provider.8=com.ibm.crypto.pkcs11.provider.IBMPKCS11
```

Implementing tokens for security attribute propagation

As part of an extensible architecture, WebSphere Application Server enables you to implement your own tokens in which to propagate security attributes.

About this task

The following topics are covered in this section:

- Implementing a custom propagation token
- Implementing a custom authorization token
- Implementing a custom a single sign-on token
- Implementing a custom authentication token
- Propagating a custom Java serializable object

Implementing a custom propagation token

This topic explains how you might create your own propagation token implementation, which is set on the running thread and propagated downstream.

About this task

The default propagation token usually is sufficient for propagating attributes that are not user-specific. Consider writing your own implementation if you want to accomplish one of the following tasks:

- Isolate your attributes within your own implementation.
- Serialize the information using custom serialization. You must deserialize the bytes at the target and add that information back on the thread by plugging in a custom login module into the inbound system login configurations. This task also might include encryption and decryption.

To implement a custom propagation token, you must complete the following steps:

1. Write a custom implementation of the `PropagationToken` interface. Many different methods are available for implementing the `PropagationToken` interface. However, make sure that the methods that are required by the `PropagationToken` interface and the token interface are fully implemented.

After you implement this interface, you can place it in the `app_server_root/classes` directory. Alternatively, you can place the class in any private directory. However, make sure that the WebSphere Application Server class loader can locate the class and that it is granted the appropriate permissions. You can add the Java archive (JAR) file or directory that contains this class into the `server.policy` file so that it has the required permissions for the server code.

Note: All of the token types that are defined by the propagation framework have similar interfaces. The token types are marker interfaces that implement the `com.ibm.wsspi.security.token.Token` interface. This interface defines most of the methods. If you plan to implement more than one

token type, consider creating an abstract class that implements the `com.ibm.wsspi.security.token.Token` interface. All of your token implementations, including the propagation token, might extend the abstract class and then most of the work is complete.

To see an implementation of the propagation token, see “Example: `com.ibm.wsspi.security.token.PropagationToken` implementation.”

2. Add and receive the custom propagation token during WebSphere Application Server logins. This task is typically accomplished by adding a custom login module to the various application and system login configurations. You also can add the implementation from an application. However, to deserialize the information, you need to plug in a custom login module, which is discussed in “Propagating a custom Java serializable object” on page 825. The `WSSecurityPropagationHelper` class has APIs that are used to set a propagation token on the thread and to retrieve the token from the thread to make updates. The code sample in “Example: Custom propagation token login module” on page 795 shows how to determine if the login is an initial login or a propagation login. The difference between these login types is whether the `WSTokenHolderCallback` callback contains propagation data. If the callback does not contain propagation data, initialize a new custom propagation token implementation and set it on the thread. If the callback contains propagation data, look for your specific custom propagation token `TokenHolder` instance, convert the byte array back into your custom `PropagationToken` object, and set it back on the thread. The code sample shows both instances.

You can add attributes any time your custom propagation token is added to the thread. If you add attributes between requests and the `getUniqueld` method changes, the Common Secure Interoperability Version 2 (CSIv2) client session is invalidated so that it can send the new information downstream. Adding attributes between requests can affect performance. In many cases, you want the downstream requests to receive the new propagation token information.

To add the custom propagation token to the thread, call the `WSSecurityPropagationHelper.addPropagationToken` method. This call requires the WebSphereRuntimePermission “setPropagationToken” Java 2 Security permission.

3. Add your custom login module to WebSphere Application Server system login configurations that already contain the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` login module for receiving serialized versions of your custom propagation token. You can also add this login module to any of the application logins where you might want to generate your custom propagation token on the thread during the login. Alternatively, you can generate the custom `PropagationToken` implementation from within your application. However, to deserialize it, you need to add the implementation to the system login modules.

For information on how to add your custom login module to the existing login configurations, see “Developing custom login modules for a system login configuration” on page 375.

Results

After completing these steps, you have implemented a custom `PropagationToken`.

Example: `com.ibm.wsspi.security.token.PropagationToken` implementation

Use this file to see an example of a propagation token implementation. The following sample code does not extend an abstract class, but implements the `com.ibm.wsspi.security.token.PropagationToken` interface directly. You can implement the interface directly, but it might cause you to write duplicate code. However, you might choose to implement the interface directly if considerable differences exist between how you handle the various token implementations.

For information on how to implement a custom propagation token, see “Implementing a custom propagation token” on page 789.

```
package com.ibm.websphere.security.token;

import com.ibm.websphere.security.WSSecurityException;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.*;
import com.ibm.websphere.security.WebSphereRuntimePermission;
```

```

import java.io.ByteArrayOutputStream;
import java.io.ByteArrayInputStream;
import java.io.DataOutputStream;
import java.io.DataInputStream;
import java.io.ObjectOutputStream;
import java.io.ObjectInputStream;
import java.io.OutputStream;
import java.io.InputStream;
import java.util.ArrayList;

public class CustomPropagationTokenImpl implements com.ibm.wsspi.security.
    token.PropagationToken
{
    private java.util.Hashtable hashtable = new java.util.Hashtable();
    private byte[] tokenBytes = null;
    // 2 hours in millis, by default
    private static long expire_period_in_millis = 2*60*60*1000;
    private long counter = 0;

/**
 * The constructor that is used to create initial PropagationToken instance
 */

    public CustomAbstractTokenImpl ()
    {
        // set the token version
        addAttribute("version", "1");
        // set the token expiration
        addAttribute("expiration", new Long(System.currentTimeMillis() +
            expire_period_in_millis).toString());
    }

/**
 * The constructor that is used to deserialize the token bytes received
 * during a propagation login.
 */
    public CustomAbstractTokenImpl (byte[] token_bytes)
    {
        try
        {
            hashtable = (java.util.Hashtable) com.ibm.wsspi.security.token.
                WSOpaqueTokenHelper.deserialize(token_bytes);
        }
        catch (Exception e)
        {
            e.printStackTrace();
        }
    }

/**
 * Validates the token including expiration, signature, and so on.
 * @return boolean
 */

    public boolean isValid ()
    {
        long expiration = getExpiration();

        // if you set the expiration to 0, it does not expire
        if (expiration != 0)
        {
            // return if this token is still valid
            long current_time = System.currentTimeMillis();

            boolean valid = ((current_time < expiration) ? true : false);
            System.out.println("isValid: returning " + valid);
            return valid;
        }
    }
}

```

```

    }
    else
    {
        System.out.println("isValid: returning true by default");
        return true;
    }
}

/**
 * Gets the expiration as a long type.
 * @return long
 */
public long getExpiration()
{
    // get the expiration value from the hashtable
    String[] expiration = getAttributes("expiration");

    if (expiration != null && expiration[0] != null)
    {
        // expiration is the first element (should only be one)
        System.out.println("getExpiration: returning " + expiration[0]);
        return new Long(expiration[0]).longValue();
    }

    System.out.println("getExpiration: returning 0");
    return 0;
}

/**
 * Returns if this token should be forwarded/propagated downstream.
 * @return boolean
 */
public boolean isForwardable()
{
    // You can choose whether your token gets propagated. In some cases
    // you might want the token to be local only.
    return true;
}

/**
 * Gets the principal that this token belongs to. If this token is an
 * authorization token, this principal string must match the authentication
 * token principal string or the message is rejected.
 * @return String
 */
public String getPrincipal()
{
    // It is not necessary for the PropagationToken to return a principal,
    // because it is not user-centric.
    return "";
}

/**
 * Returns the unique identifier of the token based upon information that
 * the provider considers makes it a unique token. This identifier is used
 * for caching purposes and might be used in combination with other token
 * unique IDs that are part of the same Subject.
 *
 * This method should return null if you want the accessID of the user to
 * represent its uniqueness. This is the typical scenario.
 *
 * @return String
 */
public String getUniqueID()
{
    // If you want to propagate the changes to this token, change the
    // value that this unique ID returns whenever the token is changed.

```

```

        // Otherwise, CSiv2 uses an existing session when everything else is
        // the same. This getUniqueID is checked by CSiv2 to determine the
        // session lookup.
    return counter;
}

/**
 * Gets the bytes to be sent across the wire. The information in the byte[]
 * needs to be enough to recreate the Token object at the target server.
 * @return byte[]
 */
public byte[] getBytes ()
{
    if (hashtable != null)
    {
        try
        {
            // Do this if the object is set to read-only during login commit
            // because this guarantees that no new data is set.
            if (isReadOnly() && tokenBytes == null)
                tokenBytes = com.ibm.wsspi.security.token.WSOpaqueTokenHelper.
                    serialize(hashtable);

            // You can deserialize this in the downstream login module using
            // WSOpaqueTokenHelper.deserialize()
            return tokenBytes;
        }
        catch (Exception e)
        {
            e.printStackTrace();
            return null;
        }
    }

    System.out.println("getBytes: returning null");
    return null;
}

/**
 * Gets the name of the token, which is used to identify the byte[] in the
 * protocol message.
 * @return String
 */
public String getName()
{
    return this.getClass().getName();
}

/**
 * Gets the version of the token as a short type. This code also is used
 * to identify the byte[] in the protocol message.
 * @return short
 */
public short getVersion()
{
    String[] version = getAttributes("version");

    if (version != null && version[0] != null)
        return new Short(version[0]).shortValue();

    System.out.println("getVersion: returning default of 1");
    return 1;
}

/**
 * When called, the token becomes irreversibly read-only. The implementation
 * needs to ensure that any setter methods check that this read-only flag has

```

```

* been set.
*/
public void setReadOnly()
{
    addAttribute("readonly", "true");
}

/**
 * Called internally to see if the token is readonly
 */
private boolean isReadOnly()
{
    String[] readonly = getAttributes("readonly");

    if (readonly != null && readonly[0] != null)
        return new Boolean(readonly[0]).booleanValue();

    System.out.println("isReadOnly: returning default of false");
    return false;
}

/**
 * Gets the attribute value based on the named value.
 * @param String key
 * @return String[]
 */
public String[] getAttributes(String key)
{
    ArrayList array = (ArrayList) hashtable.get(key);

    if (array != null && array.size() > 0)
    {
        return (String[]) array.toArray(new String[0]);
    }

    return null;
}

/**
 * Sets the attribute name and value pair. Returns the previous values set
 * for the key, or returns null if the value is not previously set.
 * @param String key
 * @param String value
 * @returns String[];
 */
public String[] addAttribute(String key, String value)
{
    // Gets the current value for the key
    ArrayList array = (ArrayList) hashtable.get(key);

    if (!isReadOnly())
    {
        // Increments the counter to change the uniqueID
        counter++;

        // Copies the ArrayList to a String[] as it currently exists
        String[] old_array = null;
        if (array != null && array.size() > 0)
            old_array = (String[]) array.toArray(new String[0]);

        // Allocates a new ArrayList if one was not found
        if (array == null)
            array = new ArrayList();

        // Adds the String to the current array list
        array.add(value);
    }
}

```

```

    // Adds the current ArrayList to the Hashtable
    hashtable.put(key, array);

    // Returns the old array
    return old_array;
}

return (String[]) array.toArray(new String[0]);
}

/**
 * Gets the list of all of the attribute names present in the token.
 * @return java.util.Enumeration
 */
public java.util.Enumeration getAttributeNames()
{
    return hashtable.keys();
}

/**
 * Returns a deep clone of this token. This is typically used by the session
 * logic of the CSiv2 server to create a copy of the token as it exists in the
 * session.
 * @return Object
 */
public Object clone()
{
    com.ibm.websphere.security.token.CustomPropagationTokenImpl deep_clone =
        new com.ibm.websphere.security.token.CustomPropagationTokenImpl();

    java.util.Enumeration keys = getAttributeNames();

    while (keys.hasMoreElements())
    {
        String key = (String) keys.nextElement();

        String[] list = (String[]) getAttributes(key);

        for (int i=0; i<list.length; i++)
            deep_clone.addAttribute(key, list[i]);
    }

    return deep_clone;
}
}

```

Example: Custom propagation token login module

This example shows how to determine if the login is an initial login or a propagation login.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        // (For more information on what to do during initialization, see
        // "Developing custom login modules for a system login configuration" on page 375.)
    }

    public boolean login() throws LoginException
    {
        // (For more information on what to do during login, see
        // "Developing custom login modules for a system login configuration" on page 375.)

        // Handles the WSTokenHolderCallback to see if this is an initial
        // or propagation login.
        Callback callbacks[] = new Callback[1];
    }
}

```



```

callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");

try
{
    callbackHandler.handle(callbacks);
}
catch (Exception e)
{
    // handle exception
}

// Receives the ArrayList of TokenHolder objects (the serialized tokens)
List authzTokenList = ((WSTokenHolderCallback) callbacks[0]).getTokenHolderList();

if (authzTokenList != null)
{
    // Iterates through the list looking for your custom token
    for (int i=0; i<authzTokenList.size(); i++)
    {
        TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

        // Looks for the name and version of your custom PropagationToken implementation
        if (tokenHolder.getName().equals("
            com.ibm.websphere.security.token.CustomPropagationTokenImpl") &&
            tokenHolder.getVersion() == 1)
        {
            // Passes the bytes into your custom PropagationToken constructor
            // to deserialize
            customPropToken = new
                com.ibm.websphere.security.token.CustomPropagationTokenImpl(tokenHolder.
                    getBytes());

        }
    }
}
else // This is not a propagation login. Create a new instance of
    // your PropagationToken implementation
{
    // Adds a new custom propagation token. This is an initial login
    customPropToken = new com.ibm.websphere.security.token.CustomPropagationTokenImpl();

    // Adds any initial attributes
    if (customPropToken != null)
    {
        customPropToken.addAttribute("key1", "value1");
        customPropToken.addAttribute("key1", "value2");
        customPropToken.addAttribute("key2", "value1");
        customPropToken.addAttribute("key3", "something different");
    }
}

// Note: You can add the token to the thread during commit in case
// something happens during the login.
}

public boolean commit() throws LoginException
{
    // For more information on what to do during commit, see
    // "Developing custom login modules for a system login configuration" on page 375
    if (customPropToken != null)
    {
        // Sets the propagation token on the thread
        try
        {

            System.out.println(tc, "*** ADDED MY CUSTOM PROPAGATION TOKEN TO THE THREAD ***");
            // Prints out the values in the deserialized propagation token

```

```

java.util.Enumeration keys = customPropToken.getAttributeNames();
while (keys.hasMoreElements())
{
    String key = (String) keys.nextElement();
    String[] list = (String[]) customPropToken.getAttributes(key);
    for (int k=0; k<list.length; k++)
        System.out.println("Key/Value: " + key + "/" + list[k]);
}

// This sets it on the thread using getName() + getVersion() as the key
com.ibm.wsspi.security.token.WSSecurityPropagationHelper.addPropagationToken(
    customPropToken);
}
catch (Exception e)
{
    // Handles exception
}

// Now you can verify that you have set it properly by trying to get
// it back from the thread and print the values.
try
{
    // This gets the PropagationToken from the thread using getName()
    // and getVersion() parameters.
    com.ibm.wsspi.security.token.PropagationToken tempPropagationToken =
        com.ibm.wsspi.security.token.WSSecurityPropagationHelper.getPropagationToken
            ("com.ibm.websphere.security.token.CustomPropagationTokenImpl", 1);

    if (tempPropagationToken != null)
    {
        System.out.println(tc, "*** RECEIVED MY CUSTOM PROPAGATION
            TOKEN FROM THE THREAD ***");
        // Prints out the values in the deserialized propagation token
        java.util.Enumeration keys = tempPropagationToken.getAttributeNames();
        while (keys.hasMoreElements())
        {
            String key = (String) keys.nextElement();
            String[] list = (String[]) tempPropagationToken.getAttributes(key);
            for (int k=0; k<list.length; k++)
                System.out.println("Key/Value: " + key + "/" + list[k]);
        }
    }
}
catch (Exception e)
{
    // Handles exception
}
}

// Defines your login module variables
com.ibm.wsspi.security.token.PropagationToken customPropToken = null;
}

```

Implementing a custom authorization token

This task explains how you might create your own `AuthorizationToken` implementation, which is set in the login Subject and propagated downstream.

About this task

The default `AuthorizationToken` usually is sufficient for propagating attributes that are user-specific. Consider writing your own implementation if you want to accomplish one of the following tasks:

- Isolate your attributes within your own implementation.
- Serialize the information using custom serialization. You must deserialize the bytes at the target and add that information back on the thread. This task also might include encryption and decryption.
- Affect the overall uniqueness of the Subject using the `getUniqueID()` application programming interface (API).

To implement a custom authorization token, you must complete the following steps:

1. Write a custom implementation of the `AuthorizationToken` interface. There are many different methods for implementing the `AuthorizationToken` interface. However, make sure that the methods required by the `AuthorizationToken` interface and the token interface are fully implemented.

After you implement this interface, you can place it in the `app_server_root/classes` directory. Alternatively, you can place the class in any private directory. However, make sure that the WebSphere Application Server class loader can locate the class and that it is granted the appropriate permissions. You can add the Java archive (JAR) file or directory that contains this class into the `server.policy` file so that it has the necessary permissions that are needed by the server code.

Note: All of the token types defined by the propagation framework have similar interfaces. Basically, the token types are marker interfaces that implement the `com.ibm.wsspi.security.token.Token` interface. This interface defines most of the methods. If you plan to implement more than one token type, consider creating an abstract class that implements the `com.ibm.wsspi.security.token.Token` interface. All of your token implementations, including the `AuthorizationToken`, might extend the abstract class and then most of the work is completed.

To see an implementation of `AuthorizationToken`, see “Example: `com.ibm.wsspi.security.token.AuthorizationToken` implementation” on page 799

2. Add and receive the custom `AuthorizationToken` during WebSphere Application Server logins. This task is typically accomplished by adding a custom login module to the various application and system login configurations. However, in order to deserialize the information, you must plug in a custom login module, which is discussed in “Propagating a custom Java serializable object” on page 825. After the object is instantiated in the login module, you can add the object to the Subject during the `commit()` method.

If you only want to add information to the Subject to get propagated, see “Propagating a custom Java serializable object” on page 825. If you want to ensure that the information is propagated, want to do your own custom serialization, or want to specify the uniqueness for Subject caching purposes, then consider writing your own `AuthorizationToken` implementation.

The code sample in “Example: custom `AuthorizationToken` login module” on page 804 shows how to determine if the login is an initial login or a propagation login. The difference between these login types is whether the `WSTokenHolderCallback` contains propagation data. If the callback does not contain propagation data, initialize a new custom `AuthorizationToken` implementation and set it into the Subject. If the callback contains propagation data, look for your specific custom `AuthorizationToken` `TokenHolder` instance, convert the `byte[]` back into your custom `AuthorizationToken` object, and set it back into the Subject. The code sample shows both instances.

You can make your `AuthorizationToken` read-only in the commit phase of the login module. If you do not make the token read-only, then attributes can be added within your applications.

3. Add your custom login module to WebSphere Application Server system login configurations that already contain the `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule` for receiving serialized versions of your custom authorization token.

Because this login module relies on information in the `sharedState` added by the `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule`, add this login module after `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule`. For information on how to add your custom login module to the existing login configurations, see “Developing custom login modules for a system login configuration” on page 375.

Results

After completing these steps, you have implemented a custom AuthorizationToken.

Example: com.ibm.wsspi.security.token.AuthorizationToken implementation

Use this file to see an example of a AuthorizationToken implementation. The following sample code does not extend an abstract class, but rather implements the com.ibm.wsspi.security.token.AuthorizationToken interface directly. You can implement the interface directly, but it might cause you to write duplicate code. However, you might choose to implement the interface directly if there are considerable differences between how you handle the various token implementations.

For information on how to implement a custom AuthorizationToken, see “Implementing a custom authorization token” on page 797.

```
package com.ibm.websphere.security.token;

import com.ibm.websphere.security.WSSecurityException;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.*;
import com.ibm.websphere.security.WebSphereRuntimePermission;
import java.io.ByteArrayOutputStream;
import java.io.ByteArrayInputStream;
import java.io.DataOutputStream;
import java.io.DataInputStream;
import java.io.ObjectOutputStream;
import java.io.ObjectInputStream;
import java.io.OutputStream;
import java.io.InputStream;
import java.util.ArrayList;

public class CustomAuthorizationTokenImpl implements com.ibm.wsspi.security.
    token.AuthorizationToken
{
    private java.util.Hashtable hashtable = new java.util.Hashtable();
    private byte[] tokenBytes = null;
    private static long expire_period_in_millis = 2*60*60*1000;
    // 2 hours in millis, by default

/**
 * Constructor used to create initial AuthorizationToken instance
 */

    public CustomAuthorizationTokenImpl (String principal)
    {
        // Sets the principal in the token
        addAttribute("principal", principal);
        // Sets the token version
        addAttribute("version", "1");
        // Sets the token expiration
        addAttribute("expiration", new Long(System.currentTimeMillis() +
            expire_period_in_millis).toString());
    }

/**
 * Constructor used to deserialize the token bytes received during a
 * propagation login.
 */
    public CustomAuthorizationTokenImpl (byte[] token_bytes)
    {
        try
        {
            hashtable = (java.util.Hashtable) com.ibm.wsspi.security.token.
                WSOPAQUETokenHelper.deserialize(token_bytes);
        }
        catch (Exception e)
    }
}
```

```

    {
        e.printStackTrace();
    }
}

/**
 * Validates the token including expiration, signature, and so on.
 * @return boolean
 */

public boolean isValid ()
{
    long expiration = getExpiration();

    // if you set the expiration to 0, it does not expire
    if (expiration != 0)
    {
        // return if this token is still valid
        long current_time = System.currentTimeMillis();

        boolean valid = ((current_time < expiration) ? true : false);
        System.out.println("isValid: returning " + valid);
        return valid;
    }
    else
    {
        System.out.println("isValid: returning true by default");
        return true;
    }
}

/**
 * Gets the expiration as a long.
 * @return long
 */
public long getExpiration()
{
    // Gets the expiration value from the hashtable
    String[] expiration = getAttributes("expiration");

    if (expiration != null && expiration[0] != null)
    {
        // The expiration is the first element. There should be only one expiration.
        System.out.println("getExpiration: returning " + expiration[0]);
        return new Long(expiration[0]).longValue();
    }

    System.out.println("getExpiration: returning 0");
    return 0;
}

/**
 * Returns if this token should be forwarded/propagated downstream.
 * @return boolean
 */
public boolean isForwardable()
{
    // You can choose whether your token gets propagated. In some cases,
    // you might want it to be local only.
    return true;
}

/**
 * Gets the principal that this Token belongs to. If this is an authorization token,
 * this principal string must match the authentication token principal string or the
 * message will be rejected.
 * @return String

```

```

*/
public String getPrincipal()
{
    // this might be any combination of attributes
    String[] principal = getAttributes("principal");

    if (principal != null && principal[0] != null)
    {
        return principal[0];
    }

    System.out.println("getExpiration: returning null");
    return null;
}

/**
 * Returns a unique identifier of the token based upon the information that provider
 * considers makes this a unique token. This will be used for caching purposes
 * and might be used in combination with other token unique IDs that are part of
 * the same Subject.
 *
 * This method should return null if you want the accessID of the user to represent
 * uniqueness. This is the typical scenario.
 *
 * @return String
 */
public String getUniqueID()
{
    // if you don't want to affect the cache lookup, just return NULL here.
    // return null;

    String cacheKeyForThisToken = "dynamic attributes";

    // if you do want to affect the cache lookup, return a string of
    // attributes that you want factored into the lookup.
    return cacheKeyForThisToken;
}

/**
 * Gets the bytes to be sent across the wire. The information in the byte[]
 * needs to be enough to recreate the Token object at the target server.
 * @return byte[]
 */
public byte[] getBytes ()
{
    if (hashtable != null)
    {
        try
        {
            // Do this if the object is set to read-only during login commit,
            // because this makes sure that no new data gets set.
            if (isReadOnly() && tokenBytes == null)
                tokenBytes = com.ibm.wsspi.security.token.WSOpaqueTokenHelper.
                    serialize(hashtable);

            // You can deserialize this in the downstream login module using
            // WSOpaqueTokenHelper.deserialize()
            return tokenBytes;
        }
        catch (Exception e)
        {
            e.printStackTrace();
            return null;
        }
    }
}

System.out.println("getBytes: returning null");

```

```

    return null;
}

/**
 * Gets the name of the token used to identify the byte[] in the protocol message.
 * @return String
 */
public String getName()
{
    return this.getClass().getName();
}

/**
 * Gets the version of the token as an short. This also is used to identify the
 * byte[] in the protocol message.
 * @return short
 */
public short getVersion()
{
    String[] version = getAttributes("version");

    if (version != null && version[0] != null)
        return new Short(version[0]).shortValue();

    System.out.println("getVersion: returning default of 1");
    return 1;
}

/**
 * When called, the token becomes irreversibly read-only. The implementation
 * needs to ensure that any setter methods check that this flag has been set.
 */
public void setReadOnly()
{
    addAttribute("readonly", "true");
}

/**
 * Called internally to see if the token is read-only
 */
private boolean isReadOnly()
{
    String[] readonly = getAttributes("readonly");

    if (readonly != null && readonly[0] != null)
        return new Boolean(readonly[0]).booleanValue();

    System.out.println("isReadOnly: returning default of false");
    return false;
}

/**
 * Gets the attribute value based on the named value.
 * @param String key
 * @return String[]
 */
public String[] getAttributes(String key)
{
    ArrayList array = (ArrayList) hashtable.get(key);

    if (array != null && array.size() > 0)
    {
        return (String[]) array.toArray(new String[0]);
    }

    return null;
}

```



```

/**
 * Sets the attribute name and value pair. Returns the previous values set for key,
 * or null if not previously set.
 * @param String key
 * @param String value
 * @returns String[];
 */
public String[] addAttribute(String key, String value)
{
    // Gets the current value for the key
    ArrayList array = (ArrayList) hashtable.get(key);

    if (!isReadOnly())
    {
        // Copies the ArrayList to a String[] as it currently exists
        String[] old_array = null;
        if (array != null && array.size() > 0)
            old_array = (String[]) array.toArray(new String[0]);

        // Allocates a new ArrayList if one was not found
        if (array == null)
            array = new ArrayList();

        // Adds the String to the current array list
        array.add(value);

        // Adds the current ArrayList to the Hashtable
        hashtable.put(key, array);

        // Returns the old array
        return old_array;
    }

    return (String[]) array.toArray(new String[0]);
}

/**
 * Gets the list of all attribute names present in the token.
 * @return java.util.Enumeration
 */
public java.util.Enumeration getAttributeNames()
{
    return hashtable.keys();
}

/**
 * Returns a deep copying of this token, if necessary.
 * @return Object
 */
public Object clone()
{
    com.ibm.websphere.security.token.CustomAuthorizationTokenImpl deep_clone =
        new com.ibm.websphere.security.token.CustomAuthorizationTokenImpl();

    java.util.Enumeration keys = getAttributeNames();

    while (keys.hasMoreElements())
    {
        String key = (String) keys.nextElement();

        String[] list = (String[]) getAttributes(key);

        for (int i=0; i<list.length; i++)
            deep_clone.addAttribute(key, list[i]);
    }
}

```

```

        return deep_clone;
    }
}

```

Example: custom AuthorizationToken login module

This file shows how to determine if the login is an initial login or a propagation login.

For information on what to do during initialization, login and commit, see “Developing custom login modules for a system login configuration” on page 375.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        _sharedState = sharedState;
    }

    public boolean login() throws LoginException
    {
        // Handles the WSTokenHolderCallback to see if this is an initial or
        // propagation login.
        Callback callbacks[] = new Callback[1];
        callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");

        try
        {
            callbackHandler.handle(callbacks);
        }
        catch (Exception e)
        {
            // Handles exception
        }

        // Receives the ArrayList of TokenHolder objects (the serialized tokens)
        List authzTokenList = ((WSTokenHolderCallback) callbacks[0]).getTokenHolderList();

        if (authzTokenList != null)
        {
            // Iterates through the list looking for your custom token
            for (int i=0; i
            for (int i=0; i<authzTokenList.size(); i++)
            {
                TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

                // Looks for the name and version of your custom AuthorizationToken
                // implementation
                if (tokenHolder.getName().equals("com.ibm.websphere.security.token.
                    CustomAuthorizationTokenImpl") &&
                    tokenHolder.getVersion() == 1)
                {
                    // Passes the bytes into your custom AuthorizationToken constructor
                    // to deserialize
                    customAuthzToken = new
                    com.ibm.websphere.security.token.CustomAuthorizationTokenImpl(
                        tokenHolder.getBytes());
                }
            }
        }
        else
        {
            // This is not a propagation login. Create a new instance of your
            // AuthorizationToken implementation
        {
            // Gets the principal from the default AuthenticationToken. This must match
            // all tokens.

```

```

defaultAuthToken = (com.ibm.wsspi.security.token.AuthenticationToken)
    sharedState.get(com.ibm.wsspi.security.auth.callback.Constants.WSAUTHTOKEN_KEY);
String principal = defaultAuthToken.getPrincipal();

    // Adds a new custom authorization token. This is an initial login. Pass the
    // principal into the constructor
    customAuthzToken = new com.ibm.websphere.security.token.
        CustomAuthorizationTokenImpl(principal);

// Adds any initial attributes
if (customAuthzToken != null)
{
    customAuthzToken.addAttribute("key1", "value1");
    customAuthzToken.addAttribute("key1", "value2");
    customAuthzToken.addAttribute("key2", "value1");
    customAuthzToken.addAttribute("key3", "something different");
}
}

    // Note: You can add the token to the Subject during commit in case something
    // happens during the login.
}

public boolean commit() throws LoginException
{
    if (customAut // (hzToken != null)
    {
        // sSets the customAuthzToken token into the Subject
        try
        {
            public final AuthorizationToken customAuthzTokenPriv = customAuthzToken;
            // Do this in a doPrivileged code block so that application code does not
            // need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Adds the custom authorization token if it is not null
                        // and not already in the Subject
                        if ((customAuthzTokenPriv != null) &&
                            (!subject.getPrivateCredentials().contains(customAuthzTokenPriv)))
                        {
                            subject.getPrivateCredentials().add(customAuthzTokenPriv);
                        }
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }
                }
            });
            return null;
        }
    }
}
catch (Exception e)
{
    throw new WSLoginFailedException (e.getMessage(), e);
}
}

// Defines your login module variables

```

```
com.ibm.wsspi.security.token.AuthorizationToken customAuthzToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthToken = null;
java.util.Map _sharedState = null;
}
```

Implementing a custom single sign-on token

You can create your own single sign-on token implementation. The single sign-on token implementation is set in the login Subject and added to the HTTP response as an HTTP cookie.

About this task

The cookie name is the concatenation of the `SingleSignonToken.getName` application programming interface (API) and the `SingleSignonToken.getVersion` API. There is no delimiter. When you add a single sign-on token to the Subject, it also gets propagated horizontally and downstream in case the Subject is used for other Web requests. You must deserialize your custom single sign-on token when you receive it from a propagation login. Consider writing your own implementation if you want to accomplish one of the following tasks:

- Isolate your attributes within your own implementation.
- Serialize the information using custom serialization. Encrypt the information because it is out to the HTTP response and is available on the Internet. You must deserialize or decrypt the bytes at the target and add that information back into the Subject.
- Affect the overall uniqueness of the Subject using the `getUniqueID` API.

To implement a custom single sign-on token, complete the following steps:

1. Write a custom implementation of the `SingleSignonToken` interface.

Many different methods are available for implementing the `SingleSignonToken` interface. However, make sure the methods that are required by the `SingleSignonToken` interface and the token interface are fully implemented.

After you implement this interface, you can place it in the `app_server_root/classes` directory. Alternatively, you can place the class in any private directory. However, make sure that the WebSphere Application Server class loader can locate the class and that it is granted the appropriate permissions. You can add the Java archive (JAR) file or directory that contains this class into the `server.policy` file so that it has the required permissions for the server code.

Note: All of the token types that are defined by the propagation framework have similar interfaces. Basically, the token types are marker interfaces that implement the `com.ibm.wsspi.security.token.Token` interface. This interface defines most of the methods. If you plan to implement more than one token type, consider creating an abstract class that implements the `com.ibm.wsspi.security.token.Token` interface. All of your token implementations, including the single sign-on token, might extend the abstract class and then most of the work is complete.

To see an implementation of the single sign-on token, see “Example: A `com.ibm.wsspi.security.token.SingleSignonToken` implementation” on page 807

2. Add and receive the custom single sign-on token during WebSphere Application Server logins. This task is typically accomplished by adding a custom login module to the various application and system login configurations. However, to deserialize the information, you need to plug in a custom login module, which is discussed in a subsequent step. After the object is instantiated in the login module, you can add it to the Subject during the commit method.

The code sample in “Example: A custom single sign-on token login module” on page 812, shows how to determine if the login is an initial login or a propagation login. The difference is whether the `WSTokenHolderCallback` callback contains propagation data. If the callback does not contain propagation data, initialize a new custom single sign-on token implementation and set it into the Subject. Also, look for the HTTP cookie from the HTTP request if the HTTP request object is available in the callback. You can get your custom single sign-on token both from a horizontal propagation login

and from the HTTP request. However, it is recommended that you make the token available in both places because then the information arrives at any front-end application server, even if that server does not support propagation.

You can make your single sign-on token read-only in the commit phase of the login module. If you make the token read-only, additional attributes cannot be added within your applications.

Note:

- HTTP cookies have a size limitation. Size restrictions should be included in the documentation for your specific browser.
 - The WebSphere Application Server runtime does not handle cookies that it does not generate, so this cookie is not used by the runtime.
 - The SingleSignonToken object, when in the Subject, does affect the cache lookup of the Subject if you return something in the getUniqueID method.
3. Get the HTTP cookie from the HTTP request object during login or from an application. The sample code that is found in “Example: An HTTP cookie retrieval” on page 814 shows how you can retrieve the HTTP cookie from the HTTP request, decode the cookie so that it is back to your original bytes, and create your custom SingleSignonToken object from the bytes.
 4. Add your custom login module to WebSphere Application Server system login configurations that already contain the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` for receiving serialized versions of your custom propagation token. Because this login module relies on information in the `sharedState` state that is added by the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` login module, add this login module after the `com.ibm.ws.security.server.lm.wsMapDefaultInboundLoginModule` login module.
For information on adding your custom login module into the existing login configurations, see “Developing custom login modules for a system login configuration” on page 375.

Results

After completing these steps, you have implemented a custom single sign-on token.

Example: A `com.ibm.wsspi.security.token.SingleSignonToken` implementation

Use this file to see an example of a single sign-on implementation. The following sample code does not extend an abstract class, but implements the `com.ibm.wsspi.security.token.SingleSignonToken` interface directly. You can implement the interface directly, but it might cause you to write duplicate code. However, you might choose to implement the interface directly if considerable differences exist between how you handle the various token implementations.

For information on how to implement a custom single sign-on token, see “Implementing a custom single sign-on token” on page 806.

```
package com.ibm.websphere.security.token;

import com.ibm.websphere.security.WSSecurityException;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.*;
import com.ibm.websphere.security.WebSphereRuntimePermission;
import java.io.ByteArrayOutputStream;
import java.io.ByteArrayInputStream;
import java.io.DataOutputStream;
import java.io.DataInputStream;
import java.io.ObjectOutputStream;
import java.io.ObjectInputStream;
import java.io.OutputStream;
import java.io.InputStream;
import java.util.ArrayList;

public class CustomSingleSignonTokenImpl implements com.ibm.wsspi.security.
    token.SingleSignonToken
```

```

{
private java.util.Hashtable hashtable = new java.util.Hashtable();
private byte[] tokenBytes = null;
    // 2 hours in millis, by default
private static long expire_period_in_millis = 2*60*60*1000;

/**
 * Constructor used to create initial SingleSignonToken instance
 */

public CustomSingleSignonTokenImpl (String principal)
{
    // set the principal in the token
    addAttribute("principal", principal);
    // set the token version
    addAttribute("version", "1");
    // set the token expiration
    addAttribute("expiration", new Long(System.currentTimeMillis() +
        expire_period_in_millis).toString());
}

/**
 * Constructor used to deserialize the token bytes received during a propagation login.
 */
public CustomSingleSignonTokenImpl (byte[] token_bytes)
{
    try
    {
        // you should implement a decryption algorithm to decrypt the cookie bytes
        hashtable = (java.util.Hashtable) some_decryption_algorithm (token_bytes);
    }
    catch (Exception e)
    {
        e.printStackTrace();
    }
}

/**
 * Validates the token including expiration, signature, and so on.
 * @return boolean
 */

public boolean isValid ()
{
    long expiration = getExpiration();

    // if you set the expiration to 0, it does not expire
    if (expiration != 0)
    {
        // return if this token is still valid
        long current_time = System.currentTimeMillis();

        boolean valid = ((current_time < expiration) ? true : false);
        System.out.println("isValid: returning " + valid);
        return valid;
    }
    else
    {
        System.out.println("isValid: returning true by default");
        return true;
    }
}

/**
 * Gets the expiration as a long.
 * @return long
 */

```

```

public long getExpiration()
{
    // get the expiration value from the hashtable
    String[] expiration = getAttributes("expiration");

    if (expiration != null && expiration[0] != null)
    {
        // expiration will always be the first element (should only be one)
        System.out.println("getExpiration: returning " + expiration[0]);
        return new Long(expiration[0]).longValue();
    }

    System.out.println("getExpiration: returning 0");
    return 0;
}

/**
 * Returns if this token should be forwarded/propagated downstream.
 * @return boolean
 */
public boolean isForwardable()
{
    // You can choose whether your token gets propagated or not, in some cases
    // you might want it to be local only.
    return true;
}

/**
 * Gets the principal that this Token belongs to. If this is an authorization token,
 * this principal string must match the authentication token principal string or the
 * message will be rejected.
 * @return String
 */
public String getPrincipal()
{
    // this could be any combination of attributes
    String[] principal = getAttributes("principal");

    if (principal != null && principal[0] != null)
    {
        return principal[0];
    }

    System.out.println("getExpiration: returning null");
    return null;
}

/**
 * Returns a unique identifier of the token based upon information the provider
 * considers makes this a unique token. This will be used for caching purposes
 * and may be used in combination with other token unique IDs that are part of
 * the same Subject.
 *
 * This method should return null if you want the access ID of the user to represent
 * uniqueness. This is the typical scenario.
 *
 * @return String
 */
public String getUniqueID()
{
    // this could be any combination of attributes
    return getPrincipal();
}

/**
 * Gets the bytes to be sent across the wire. The information in the byte[]
 * needs to be enough to recreate the Token object at the target server.

```



```

* @return byte[]
*/
public byte[] getBytes ()
{
    if (hashtable != null)
    {
        try
        {
            // do this if the object is set read-only during login commit,
            // since this guarantees no new data gets set.
            if (isReadOnly() && tokenBytes == null)
                tokenBytes = some_encryption_algorithm (hashtable);

            // you can deserialize the tokenBytes using a similiar decryption algorithm.
            return tokenBytes;
        }
        catch (Exception e)
        {
            e.printStackTrace();
            return null;
        }
    }

    System.out.println("getBytes: returning null");
    return null;
}

/**
 * Gets the name of the token, used to identify the byte[] in the protocol message.
 * @return String
 */
public String getName()
{
    return "myCookieName";
}

/**
 * Gets the version of the token as a short. This is also used to identify the
 * byte[] in the protocol message.
 * @return short
 */
public short getVersion()
{
    String[] version = getAttributes("version");

    if (version != null && version[0] != null)
        return new Short(version[0]).shortValue();

    System.out.println("getVersion: returning default of 1");
    return 1;
}

/**
 * When called, the token becomes irreversibly read-only. The implementation
 * needs to ensure any setter methods check that this has been set.
 */
public void setReadOnly()
{
    addAttribute("readonly", "true");
}

/**
 * Called internally to see if the token is readonly
 */
private boolean isReadOnly()
{
    String[] readonly = getAttributes("readonly");

```

```

    if (readonly != null && readonly[0] != null)
        return new Boolean(readonly[0]).booleanValue();

    System.out.println("isReadOnly: returning default of false");
    return false;
}

/**
 * Gets the attribute value based on the named value.
 * @param String key
 * @return String[]
 */
public String[] getAttributes(String key)
{
    ArrayList array = (ArrayList) hashtable.get(key);

    if (array != null && array.size() > 0)
    {
        return (String[]) array.toArray(new String[0]);
    }

    return null;
}

/**
 * Sets the attribute name/value pair. Returns the previous values set for key,
 * or null if not previously set.
 * @param String key
 * @param String value
 * @returns String[];
 */
public String[] addAttribute(String key, String value)
{
    // get the current value for the key
    ArrayList array = (ArrayList) hashtable.get(key);

    if (!isReadOnly())
    {
        // copy the ArrayList to a String[] as it currently exists
        String[] old_array = null;
        if (array != null && array.size() > 0)
            old_array = (String[]) array.toArray(new String[0]);

        // allocate a new ArrayList if one was not found
        if (array == null)
            array = new ArrayList();

        // add the String to the current array list
        array.add(value);

        // add the current ArrayList to the Hashtable
        hashtable.put(key, array);

        // return the old array
        return old_array;
    }

    return (String[]) array.toArray(new String[0]);
}

/**
 * Gets the List of all attribute names present in the token.
 * @return java.util.Enumeration
 */
public java.util.Enumeration getAttributeNames()

```

```

{
    return hashtable.keys();
}

/**
 * Returns a deep copying of this token, if necessary.
 * @return Object
 */
public Object clone()
{
    com.ibm.websphere.security.token.CustomSingleSignonImpl deep_clone =
        new com.ibm.websphere.security.token.CustomSingleSignonTokenImpl();

    java.util.Enumeration keys = getAttributeNames();

    while (keys.hasMoreElements())
    {
        String key = (String) keys.nextElement();

        String[] list = (String[]) getAttributes(key);

        for (int i=0; i<list.length; i++)
            deep_clone.addAttribute(key, list[i]);
    }

    return deep_clone;
}
}

```

Example: A custom single sign-on token login module

This file shows how to determine if the login is an initial login or a propagation login.

For information on initialization and on what to do during login and commit, see “Developing custom login modules for a system login configuration” on page 375.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        _sharedState = sharedState;
    }

    public boolean login() throws LoginException
    {
        // Handles the WSTokenHolderCallback to see if this is an initial or
        // propagation login.
        Callback callbacks[] = new Callback[1];
        callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");

        try
        {
            callbackHandler.handle(callbacks);
        }
        catch (Exception e)
        {
            // handle exception
        }

        // Receives the ArrayList of TokenHolder objects (the serialized tokens)
        List authzTokenList = ((WSTokenHolderCallback) callbacks[0]).getTokenHolderList();

        if (authzTokenList != null)
        {
            // iterate through the list looking for your custom token
            for (int i=0; i
            for (int i=0; i<authzTokenList.size(); i++)

```

```

{
    TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

    // Looks for the name and version of your custom SingleSignonToken
    // implementation
    if (tokenHolder.getName().equals("myCookieName")
        && tokenHolder.getVersion() == 1)
    {
        // Passes the bytes into your custom SingleSignonToken constructor
        // to deserialize
        customSSOToken = new
            com.ibm.websphere.security.token.CustomSingleSignonTokenImpl
                (tokenHolder.getBytes());
    }
}
}
else
    // This is not a propagation login. Create a new instance of your
    // SingleSignonToken implementation
    {
        // Gets the principal from the default SingleSignonToken. This principal
        // must match all tokens.
        defaultAuthToken = (com.ibm.wsspi.security.token.AuthenticationToken)
            sharedState.get(com.ibm.wsspi.security.auth.callback.Constants.WSAUTHTOKEN_KEY);
        String principal = defaultAuthToken.getPrincipal();

        // Adds a new custom single sign-on (SSO) token. This is an initial login.
        // Pass the principal into the constructor
        customSSOToken = new com.ibm.websphere.security.token.
            CustomSingleSignonTokenImpl(principal);

        // add any initial attributes
        if (customSSOToken != null)
        {
            customSSOToken.addAttribute("key1", "value1");
            customSSOToken.addAttribute("key1", "value2");
            customSSOToken.addAttribute("key2", "value1");
            customSSOToken.addAttribute("key3", "something different");
        }
    }

    // Note: You can add the token to the Subject during commit in case something
    // happens during the login.
}

public boolean commit() throws LoginException
{
    if (customSSOToken != null)
    {
        // Sets the customSSOToken token into the Subject
        try
        {
            public final SingleSignonToken customSSOTokenPriv = customSSOToken;
            // Do this in a doPrivileged code block so that application code does not
            // need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    {
                        try
                        {
                            // Adds the custom SSO token if it is not null and
                            // not already in the Subject
                            if ((customSSOTokenPriv != null) &&
                                (!subject.getPrivateCredentials().
                                    contains(customSSOTokenPriv)))

```

```

        {
            subject.getPrivateCredentials().
                add(customSSOTokenPriv);
        }
    }
    catch (Exception e)
    {
        throw new WSLoginFailedException (e.getMessage(), e);
    }

    return null;
}
});
}
catch (Exception e)
{
    throw new WSLoginFailedException (e.getMessage(), e);
}
}
}

// Defines your login module variables
com.ibm.wsspi.security.token.SingleSignonToken customSSOToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthToken = null;
java.util.Map _sharedState = null;
}

```

Example: An HTTP cookie retrieval

The following example shows you how to retrieve a cookie from an HTTP request, decode the cookie so that it is back to your original bytes, and create your custom SingleSignonToken object from the bytes. This example shows how to complete these steps from a login module. However, you also can complete these steps using a servlet.

For information on what to do during initialization, login and commit, see “Developing custom login modules for a system login configuration” on page 375.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        _sharedState = sharedState;
    }

    public boolean login() throws LoginException
    {
        // Handles the WSTokenHolderCallback to see if this is an
        // initial or propagation login.
        Callback callbacks[] = new Callback[2];
        callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");
        callbacks[1] = new WSServletRequestCallback("HttpServletRequest: ");

        try
        {
            callbackHandler.handle(callbacks);
        }
        catch (Exception e)
        {
            // Handles the exception
        }

        // receive the ArrayList of TokenHolder objects (the serialized tokens)
        List authzTokenList = ((WSTokenHolderCallback) callbacks[0]).getTokenHolderList();
        javax.servlet.http.HttpServletRequest request =
            ((WSServletRequestCallback) callbacks[1]).getHttpServletRequest();
    }
}

```

```

if (request != null)
{
    // Checks if the cookie is present
    javax.servlet.http.Cookie[] cookies = request.getCookies();
    String[] cookieStrings = getCookieValues (cookies, "myCookieName1");

    if (cookieStrings != null)
    {
        String cookieVal = null;
        for (int n=0;n<cookieStrings.length;n++)
        {
            cookieVal = cookieStrings[n];
            if (cookieVal.length()>0)
            {
                // Removes the cookie encoding from the cookie to get
                // your custom bytes
                byte[] cookieBytes =
                    com.ibm.websphere.security.WSSecurityHelper.
                        convertCookieStringToBytes(cookieVal);
                customSSOToken =
                    new com.ibm.websphere.security.token.
                        CustomSingleSignonTokenImpl(cookieBytes);

                // Now that you have your cookie from the request,
                // you can do something with it here, or add it
                // to the Subject in the commit() method for use later.
                if (debug || tc.isDebugEnabled())
                {
                    System.out.println("*** GOT MY CUSTOM SSO TOKEN FROM
                        THE REQUEST ***");
                }
            }
        }
    }
}

public boolean commit() throws LoginException
{
    if (customSSOToken != null)
    {
        // Sets the customSSOToken token into the Subject
        try
        {
            public final SingleSignonToken customSSOTokenPriv = customSSOToken;
            // Do this in a doPrivileged code block so that application code does not
            // need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Add the custom SSO token if it is not null and not
                        // already in the Subject
                        if ((customSSOTokenPriv != null) &&
                            (!subject.getPrivateCredentials().
                                contains(customSSOTokenPriv)))
                        {
                            subject.getPrivateCredentials().add(customSSOTokenPriv);
                        }
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }
                }
            });
        }
    }
}

```

```

    }
    return null;
  }
});
}
catch (Exception e)
{
  throw new WSLoginFailedException (e.getMessage(), e);
}
}
}

// Private method to get the specific cookie from the request
private String[] getCookieValues (Cookie[] cookies, String hdrName)
{
  Vector retValues = new Vector();
  int numMatches=0;
  if (cookies != null)
  {
    for (int i = 0; i < cookies.length; ++i)
    {
      if (hdrName.equals(cookies[i].getName()))
      {
        retValues.add(cookies[i].getValue());
        numMatches++;
        System.out.println(cookies[i].getValue());
      }
    }
  }

  if (retValues.size(>0)
    return (String[]) retValues.toArray(new String[numMatches]);
  else
    return null;
}

// Defines your login module variables
com.ibm.wsspi.security.token.SingleSignonToken customSSOToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthToken = null;
java.util.Map _sharedState = null;
}

```

Implementing a custom authentication token

This topic explains how you might create your own authentication token implementation, which is set in the login Subject and propagated downstream.

About this task

With this implementation you can specify an authentication token that can be used by a custom login module or application. Consider writing your own implementation if you want to accomplish one of the following tasks:

- Isolate your attributes within your own implementation.
- Serialize the information using custom serialization. You must deserialize the bytes at the target and add that information back on the thread. This task also might include encryption and decryption.
- Affect the overall uniqueness of the Subject using the `getUniqueID` application programming interface (API).

Note: Custom authentication token implementations are not used by the security runtime in WebSphere Application Server to enforce authentication. WebSphere Application Security runtime uses this token in the following situations only:

- Call the `getBytes` method for serialization
- Call the `getForwardable` method to determine whether to serialize the authentication token.
- Call the `getUniqueld` method for uniqueness
- Call the `getName` and the `getVersion` methods for adding serialized bytes to the token holder that is sent downstream

All of the other uses are custom implementations.

To implement a custom authentication token, you must complete the following steps:

1. Write a custom implementation of the `AuthenticationToken` interface. Many different methods are available for implementing the `AuthenticationToken` interface. However, make sure the methods that are required by the `AuthenticationToken` interface and the token interface are fully implemented. After you implement this interface, you can place it in the `install_dir/classes` directory. Alternatively, you can place the class in any private directory. However, make sure that the WebSphere Application Server class loader can locate the class and that it is granted the appropriate permissions. You can add the Java archive (JAR) file or directory that contains this class into the `server.policy` file so the class has the necessary permissions required by the server code.

Note: All of the token types that are defined by the propagation framework have similar interfaces. The token types are marker interfaces that implement the `com.ibm.wsspi.security.token.Token` interface. This interface defines most of the methods. If you plan to implement more than one token type, consider creating an abstract class that implements the `com.ibm.wsspi.security.token.Token` interface. All of your token implementations, including the authentication token, might extend the abstract class and then most of the work is complete.

To see an implementation of the `AuthenticationToken` interface, see “Example: A `com.ibm.wsspi.security.token.AuthenticationToken` implementation” on page 818.

2. Add and receive the custom authentication token during WebSphere Application Server logins. This task is typically accomplished by adding a custom login module to the various application and system login configurations. However, to deserialize the information you must plug in a custom login module. After the object is instantiated in the login module, you can add the object to the `Subject` during the `commit` method.

If you only want to add information to the `Subject` to get propagated, see “Propagating a custom Java serializable object” on page 825. If you want to ensure that the information is propagated, do your own custom serialization, or specify the uniqueness for `Subject` caching purposes, consider writing your own authentication token implementation.

The code sample in “Example: A custom authentication token login module” on page 823, shows how to determine if the login is an initial login or a propagation login. The difference between these login types is whether the `WSTokenHolderCallback` callback contains propagation data. If the callback does not contain propagation data, initialize a new custom authentication token implementation and set it into the `Subject`. If the callback contains propagation data, look for your specific custom authentication token `TokenHolder` instance, convert the byte array back into your custom `AuthenticationToken` object, and set it back into the `Subject`. The code sample shows both instances.

You can make your authentication token read-only in the `commit` phase of the login module. If you do not make the token read-only, attributes can be added within your applications.

3. Add your custom login module to WebSphere Application Server system login configurations that already contain the `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule` login module for receiving serialized versions of your custom authorization token.

Because this login module relies on information in the shared state that is added by the `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule` login module, add this login module after the `com.ibm.ws.security.server.Im.wsMapDefaultInboundLoginModule` login module. For information on how to add your custom login module to the existing login configurations, see “Developing custom login modules for a system login configuration” on page 375.

Results

After completing these steps, you have implemented a custom authentication token.

Example: A `com.ibm.wsspi.security.token.AuthenticationToken` implementation

The following example illustrates an authentication token implementation. The following sample code does not extend an abstract class, but rather implements the `com.ibm.wsspi.security.token.AuthenticationToken` interface directly. You can implement the interface directly, but it might cause you to write duplicate code. However, you might choose to implement the interface directly if considerable differences exist between how you handle the various token implementations.

```
package com.ibm.websphere.security.token;

import com.ibm.websphere.security.WSSecurityException;
import com.ibm.websphere.security.auth.WSLoginFailedException;
import com.ibm.wsspi.security.token.*;
import com.ibm.websphere.security.WebSphereRuntimePermission;
import java.io.ByteArrayOutputStream;
import java.io.ByteArrayInputStream;
import java.io.DataOutputStream;
import java.io.DataInputStream;
import java.io.ObjectOutputStream;
import java.io.ObjectInputStream;
import java.io.OutputStream;
import java.io.InputStream;
import java.util.ArrayList;

public class CustomAuthenticationTokenImpl implements com.ibm.wsspi.security.
    token.AuthenticationToken
{
    private java.util.Hashtable hashtable = new java.util.Hashtable();
    private byte[] tokenBytes = null;
    // 2 hours in millis, by default
    private static long expire_period_in_millis = 2*60*60*1000;
    private String oidName = "your_oid_name";
    // This string can really be anything if you do not want to use an OID.

/**
 * Constructor used to create initial AuthenticationToken instance
 */
    public CustomAuthenticationTokenImpl (String principal)
    {
        // Sets the principal in the token
        addAttribute("principal", principal);
        // Sets the token version
        addAttribute("version", "1");
        // Sets the token expiration
        addAttribute("expiration", new Long(System.currentTimeMillis()
            + expire_period_in_millis).toString());
    }

/**
 * Constructor used to deserialize the token bytes received during a
 * propagation login.
 */
    public CustomAuthenticationTokenImpl (byte[] token_bytes)
    {
        try
        {
            // The data in token_bytes should be signed and encrypted if the
            // hashtable is acting as an authentication token.
            hashtable = (java.util.Hashtable) custom_decryption_algorithm (token_bytes);
        }
        catch (Exception e)
        {
            e.printStackTrace();
        }
    }
}
```

```

    }
}

/**
 * Validates the token including expiration, signature, and so on.
 * @return boolean
 */

public boolean isValid ()
{
    long expiration = getExpiration();

    // If you set the expiration to 0, the token does not expire
    if (expiration != 0)
    {
        // Returns a response that identifies whether this token is still valid
        long current_time = System.currentTimeMillis();

        boolean valid = ((current_time < expiration) ? true : false);
        System.out.println("isValid: returning " + valid);
        return valid;
    }
    else
    {
        System.out.println("isValid: returning true by default");
        return true;
    }
}

/**
 * Gets the expiration as a long type.
 * @return long
 */
public long getExpiration()
{
    // Gets the expiration value from the hashtable
    String[] expiration = getAttributes("expiration");

    if (expiration != null && expiration[0] != null)
    {
        // The expiration is the first element and there should only be one expiration
        System.out.println("getExpiration: returning " + expiration[0]);
        return new Long(expiration[0]).longValue();
    }

    System.out.println("getExpiration: returning 0");
    return 0;
}

/**
 * Returns if this token should be forwarded/propagated downstream.
 * @return boolean
 */
public boolean isForwardable()
{
    // You can choose whether your token gets propagated. In some cases
    // you might want it to be local only.
    return true;
}

/**
 * Gets the principal to which this token belongs. If this is an
 * authorization token, this principal string must match the
 * authentication token principal string or the message is rejected.
 * @return String
 */
public String getPrincipal()

```

```

{
// This value might be any combination of attributes
String[] principal = getAttributes("principal");

if (principal != null && principal[0] != null)
{
return principal[0];
}

System.out.println("getExpiration: returning null");
return null;
}

/**
 * Returns a unique identifier of the token based upon information the provider
 * considers makes this a unique token. This identifier is used for caching purposes
 * and can be used in combination with other token unique IDs that are part of
 * the same Subject.
 *
 * This method should return null if you want the accessID of the user to represent
 * uniqueness. This is the typical scenario.
 *
 * @return String
 */
public String getUniqueID()
{
// If you do not want to affect the cache lookup, just return NULL here.
return null;

String cacheKeyForThisToken = "dynamic attributes";

// If you do want to affect the cache lookup, return a string of
// attributes that you want factored into the lookup.
return cacheKeyForThisToken;
}

/**
 * Gets the bytes to be sent across the wire. The information in the byte[]
 * needs to be enough to recreate the token object at the target server.
 * @return byte[]
 */
public byte[] getBytes ()
{
if (hashtable != null)
{
try
{
// Do this if the object is set read-only during login commit
// because this ensures that new data is not set.
if (isReadOnly() && tokenBytes == null)
tokenBytes = custom_encryption_algorithm (hashtable);

return tokenBytes;
}
catch (Exception e)
{
e.printStackTrace();
return null;
}
}

System.out.println("getBytes: returning null");
return null;
}

/**
 * Gets the name of the token, which is used to identify the byte[] in the

```

```

* protocol message.
* @return String
*/
public String getName()
{
    return oidName;
}

/**
 * Gets the version of the token as a short type. This also is used
 * to identify the byte[] in the protocol message.
 * @return short
 */
public short getVersion()
{
    String[] version = getAttributes("version");

    if (version != null && version[0] != null)
        return new Short(version[0]).shortValue();

    System.out.println("getVersion: returning default of 1");
    return 1;
}

/**
 * When called, the token becomes irreversibly read-only. The implementation
 * needs to ensure that any set methods check that this state has been set.
 */
public void setReadOnly()
{
    addAttribute("readonly", "true");
}

/**
 * Called internally to see if the token is read-only
 */
private boolean isReadOnly()
{
    String[] readonly = getAttributes("readonly");

    if (readonly != null && readonly[0] != null)
        return new Boolean(readonly[0]).booleanValue();

    System.out.println("isReadOnly: returning default of false");
    return false;
}

/**
 * Gets the attribute value based on the named value.
 * @param String key
 * @return String[]
 */
public String[] getAttributes(String key)
{
    ArrayList array = (ArrayList) hashtable.get(key);

    if (array != null && array.size() > 0)
    {
        return (String[]) array.toArray(new String[0]);
    }

    return null;
}

/**
 * Sets the attribute name/value pair. Returns the previous values set for key,
 * or null if not previously set.

```

```

* @param String key
* @param String value
* @returns String[];
*/
public String[] addAttribute(String key, String value)
{
    // Gets the current value for the key
    ArrayList array = (ArrayList) hashtable.get(key);

    if (!isReadOnly())
    {
        // Copies the ArrayList to a String[] as it currently exists
        String[] old_array = null;
        if (array != null && array.size() > 0)
            old_array = (String[]) array.toArray(new String[0]);

        // Allocates a new ArrayList if one was not found
        if (array == null)
            array = new ArrayList();

        // Adds the String to the current array list
        array.add(value);

        // Adds the current ArrayList to the Hashtable
        hashtable.put(key, array);

        // Returns the old array
        return old_array;
    }

    return (String[]) array.toArray(new String[0]);
}

/**
 * Gets the list of all attribute names present in the token.
 * @return java.util.Enumeration
 */
public java.util.Enumeration getAttributeNames()
{
    return hashtable.keys();
}

/**
 * Returns a deep copying of this token, if necessary.
 * @return Object
 */
public Object clone()
{
    com.ibm.wsspi.security.token.AuthenticationToken deep_clone =
        new com.ibm.websphere.security.token.CustomAuthenticationTokenImpl();

    java.util.Enumeration keys = getAttributeNames();

    while (keys.hasMoreElements())
    {
        String key = (String) keys.nextElement();

        String[] list = (String[]) getAttributes(key);

        for (int i=0; i<list.length; i++)
            deep_clone.addAttribute(key, list[i]);
    }

    return deep_clone;
}

```

```

/**
 * This method returns true if this token is storing a user ID and password
 * instead of a token.
 * @return boolean
 */
public boolean isBasicAuth()
{
    return false;
}
}

```

Example: A custom authentication token login module

This examples shows how to determine if the login is an initial login or a propagation login.

For information on what to do during initialization, login and commit, see “Developing custom login modules for a system login configuration” on page 375.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        _sharedState = sharedState;
    }

    public boolean login() throws LoginException
    {
        // Handles the WSTokenHolderCallback to see if this is an initial or
        // propagation login.
        Callback callbacks[] = new Callback[1];
        callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");

        try
        {
            callbackHandler.handle(callbacks);
        }
        catch (Exception e)
        {
            // Handles exception
        }

        // Receives the ArrayList of TokenHolder objects (the serialized tokens)
        List authzTokenList = ((WSTokenHolderCallback) callbacks[0]).getTokenHolderList();

        if (authzTokenList != null)
        {
            // Iterates through the list looking for your custom token
            for (int i=0; i<authzTokenList.size(); i++)
            {
                TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

                // Looks for the name and version of your custom AuthenticationToken
                // implementation
                if (tokenHolder.getName().equals("your_oid_name") && tokenHolder.getVersion() == 1)
                {
                    // Passes the bytes into your custom AuthenticationToken constructor
                    // to deserialize
                    customAuthzToken = new
                    com.ibm.websphere.security.token.
                    CustomAuthenticationTokenImpl(tokenHolder.getBytes());
                }
            }
        }
        else
            // This is not a propagation login. Create a new instance of your
            // AuthenticationToken implementation

```



```

{
    // Gets the principal from the default AuthenticationToken. This principal
    // should match all default tokens.
    // Note: WebSphere Application Server runtime only enforces this for
    // default tokens. Thus, you can choose
    // to do this for custom tokens, but it is not required.
    defaultAuthToken = (com.ibm.wsspi.security.token.AuthenticationToken)
        sharedState.get(com.ibm.wsspi.security.auth.callback.Constants.WSAUTHTOKEN_KEY);
    String principal = defaultAuthToken.getPrincipal();

    // Adds a new custom authentication token. This is an initial login. Pass
    // the principal into the constructor
    customAuthToken = new com.ibm.websphere.security.token.
        CustomAuthenticationTokenImpl(principal);

    // Adds any initial attributes
    if (customAuthToken != null)
    {
        customAuthToken.addAttribute("key1", "value1");
        customAuthToken.addAttribute("key1", "value2");
        customAuthToken.addAttribute("key2", "value1");
        customAuthToken.addAttribute("key3", "something different");
    }
}

// Note: You can add the token to the Subject during commit in case
// something happens during the login.
}

public boolean commit() throws LoginException
{
    if (customAuthToken != null)
    {
        // Sets the customAuthToken token into the Subject
        try
        {
            private final AuthenticationToken customAuthTokenPriv = customAuthToken;
            // Do this in a doPrivileged code block so that application code does
            // not need to add additional permissions
            java.security.AccessController.doPrivileged(new java.security.PrivilegedAction()
            {
                public Object run()
                {
                    try
                    {
                        // Adds the custom Authentication token if it is not
                        // null and not already in the Subject
                        if ((customAuthTokenPriv != null) &&
                            (!subject.getPrivateCredentials().
                                contains(customAuthTokenPriv)))
                        {
                            subject.getPrivateCredentials().add(customAuthTokenPriv);
                        }
                    }
                    catch (Exception e)
                    {
                        throw new WSLoginFailedException (e.getMessage(), e);
                    }
                }
            });
        }
        catch (Exception e)
        {
            throw new WSLoginFailedException (e.getMessage(), e);
        }
    }
}

```

```

    }
}

// Defines your login module variables
com.ibm.wsspi.security.token.AuthenticationToken customAuthToken = null;
com.ibm.wsspi.security.token.AuthenticationToken defaultAuthToken = null;
java.util.Map _sharedState = null;
}

```

Propagating a custom Java serializable object

This document describes how to add an object into the Subject from a login module and describes other infrastructure considerations to make sure that the Java object gets propagated.

Before you begin

Prior to completing this task, verify that security propagation is enabled in the administrative console.

About this task

With security attribute propagation enabled, you can propagate data either horizontally with single sign-on (SSO) enabled or downstream using Common Secure Interoperability Version 2 (CSIv2). When a login occurs, either through an application login configuration or a system login configuration, a custom login module can be plugged in to add Java serialized objects into the Subject during login. This document describes how to add an object into the Subject from a login module and describes other infrastructure considerations to make sure that the Java object gets propagated.

1. Add your custom Java object into the Subject from a custom login module. A two-phase process exists for each Java Authentication and Authorization Service (JAAS) login module. WebSphere Application Server completes the following processes for each login module present in the configuration:

login method

In this step, the login configuration callbacks are analyzed, if necessary, and the new objects or credentials are created.

commit method

In this step, the objects or credentials that are created during login are added into the Subject.

After a custom Java object is added into the Subject, WebSphere Application Server serializes the object on the sending server, deserializes the object on the receiving server, and adds the object back into the Subject downstream. However, some requirements exist for this process to occur successfully. For more information on the JAAS programming model, see the JAAS information provided in Security: Resources for learning.

Note: Whenever you plug a custom login module into the login infrastructure of WebSphere Application Server, make sure that the code is trusted. When you put the classes together in a Java archive (JAR) file and add the file to the `app_server_root/lib/ext/` directory, the login module has Java 2 Security AllPermissions permissions. It is recommended that you add your login module and other infrastructure classes into any private directory. However, you must modify the `profile_root/properties/server.policy` file to make sure that your private directory, Java archive (JAR) file, or both have the permissions required to run the application programming interfaces (API) that are called from the login module. Because the login module might be run after the application code on the call stack, you might add `doPrivileged` code so that you do not need to add additional properties to your applications.

The following code sample shows how to add `doPrivileged` code. For information on what to do during initialization, login and commit, see “Developing custom login modules for a system login configuration” on page 375.

```

public customLoginModule()
{
    public void initialize(Subject subject, CallbackHandler callbackHandler,

```

```

    Map sharedState, Map options)
{
}

public boolean login() throws LoginException
{
    // Construct callback for the WSTokenHolderCallback so that you
    // can determine if
    // your custom object has propagated
    Callback callbacks[] = new Callback[1];
    callbacks[0] = new WSTokenHolderCallback("Authz Token List: ");

    try
    {
        _callbackHandler.handle(callbacks);
    }
    catch (Exception e)
    {
        throw new LoginException (e.getLocalizedName());
    }

    // Checks to see if any information is propagated into this login
    List authzTokenList = ((WSTokenHolderCallback) callbacks[1]).
        getTokenHolderList();

    if (authzTokenList != null)
    {
        for (int i = 0; i < authzTokenList.size(); i++)
        {
            TokenHolder tokenHolder = (TokenHolder)authzTokenList.get(i);

            // Look for your custom object. Make sure you use
            // "startsWith"because there is some data appended
            // to the end of the name indicating in which Subject
            // Set it belongs. Example from getName():
            // "com.acme.CustomObject (1)". The class name is
            // generated at the sending side by calling the
            // object.getClass().getName() method. If this object
            // is deserialized by WebSphere Application Server,
            // then return it and you do not need to add it here.
            // Otherwise, you can add it below.
            // Note: If your class appears in this list and does
            // not use custom serialization (for example, an
            // implementation of the Token interface described in
            // the Propagation Token Framework), then WebSphere
            // Application Server automatically deserializes the
            // Java object for you. You might just return here if
            // it is found in the list.

            if (tokenHolder.getName().startsWith("com.acme.CustomObject"))
                return true;
        }

        // If you get to this point, then your custom object has not propagated
        myCustomObject = new com.acme.CustomObject();
        myCustomObject.put("mykey", "mydata");
    }
}

public boolean commit() throws LoginException
{
    try
    {
        // Assigns a reference to a final variable so it can be used in
        // the doPrivileged block
        final com.acme.CustomObject myCustomObjectFinal = myCustomObject;
        // Prevents your applications from needing a JAAS getPrivateCredential

```

```

        // permission.
        java.security.AccessController.doPrivileged(new java.security.
            PrivilegedExceptionAction()
        {
            public Object run() throws java.lang.Exception
            {
                // Try not to add a null object to the Subject or an object
                // that already exists.
                if (myCustomObjectFinal != null && !subject.getPrivateCredentials().
                    contains(myCustomObjectFinal))
                {
                    // This call requires a special Java 2 Security permission,
                    // see the JAAS application programming interface (API)
                    // documentation.
                    subject.getPrivateCredentials().add(myCustomObjectFinal);
                }
                return null;
            }
        });
    }
    catch (java.security.PrivilegedActionException e)
    {
        // Wraps the exception in a WSSLoginFailedException
        java.lang.Throwable myException = e.getException();
        throw new WSSLoginFailedException (myException.getMessage(), myException);
    }
}

// Defines your login module variables
com.acme.CustomObject myCustomObject = null;
}

```

2. Verify that your custom Java class implements the `java.io.Serializable` interface. An object that is added to the Subject must be serialized if you want the object to propagate. For example, the object must implement the `java.io.Serializable` interface. If the object is not serialized, the request does not fail, but the object does not propagate. To make sure an object that is added to the Subject is propagated, implement one of the token interfaces that is defined in “Security attribute propagation” on page 448 or add attributes to one of the following existing default token implementations:

AuthorizationToken

Add attributes if they are user-specific. For more information, see “Using the default authorization token” on page 456.

PropagationToken

Add attributes that are specific to an invocation. For more information, see “Using the default propagation token” on page 459.

If you are careful adding custom objects and follow all the steps to make sure that WebSphere Application Server can serialize and deserialize the object at each hop, then it is sufficient to use custom Java objects only.

3. Verify that your custom Java class exists on all of the systems that might receive the request. When you add a custom object into the Subject and expect WebSphere Application Server to propagate the object, put the class definitions together in a Java archive (JAR) file and add the file to the `app_server_root/lib/ext/` directory on all of the nodes where serialization or deserialization might occur. Also, verify that the Java class versions are the same.
4. Verify that your custom login module is configured in all of the login configurations used in your environment where you need to add your custom object during a login. Any login configuration that interacts with WebSphere Application Server generates a Subject that might be propagated outbound for an Enterprise JavaBeans (EJB) request. If you want WebSphere Application Server to propagate a custom object in all cases, make sure that the custom login module is added to every login configuration that is used in your environment. For more information, see “Developing custom login modules for a system login configuration” on page 375.

5. Verify that security attribute propagation is enabled on all of the downstream servers that receive the propagated information. When an EJB request is sent to a downstream server and security attribute propagation is disabled on that server, only the authentication token is sent for backwards compatibility. Therefore, you must review the configuration to verify that propagation is enabled in all of the cells that might receive requests. You must check several places in the administrative console to make sure propagation is fully enabled. For more information, see “Propagating security attributes among application servers” on page 453.
6. Add any custom objects to the propagation exclude list that you do not want to propagate. You can configure a property to exclude the propagation of objects that match specific class names, package names, or both. For example, you can have a custom object that is related to a specific process. If the object is propagated, it does not contain valid information. You must tell WebSphere Application Server not to propagate this object. Complete the following steps to specify the object in the propagation exclude list, using the administrative console:
 - a. Click **Security > Global security > Custom properties > New**.
 - b. Add `com.ibm.ws.security.propagationExcludeList` in the **Name** field.
 - c. Add the name of the custom object in the **Value** field. You can add a list of custom objects to the propagation exclude list, separated by a colon (:). For example, you might enter `com.acme.CustomLocalObject:com.acme.private.*`. You can enter a class name such as `com.acme.CustomLocalObject` or a package name such as `com.acme.private.*`. In this example, WebSphere Application Server does not propagate any class that equals `com.acme.CustomLocalObject` or begins with `com.acme.private`.

Although you can add custom objects to the propagation exclude list, you must be aware of a side effect. WebSphere Application Server stores the opaque token, or the serialized Subject contents, in a local cache for the life of the single sign-on (SSO) token. The life of the SSO token, which has a default of two hours, is configured in the SSO properties on the administrative console. The information that is added to the opaque token includes only the objects not in the exclude list.

If your authentication cache does not match your SSO token timeout, configure the authentication cache properties. See “Configuring the authentication cache” on page 466. It is recommended that you make your authentication cache timeout value equal to the SSO token timeout.

Results

As a result of this task, custom Java serializable objects are propagated horizontally or downstream. For more information on the differences between horizontal and downstream propagation, see “Security attribute propagation” on page 448.

Developing a custom interceptor for trust associations

You can define the interceptor class method that you want to use. WebSphere Application Server supports two trust association interceptor interfaces: `com.ibm.websphere.security.TrustAssociationInterceptor` and `com.ibm.wsspi.security.tai.TrustAssociationInterceptor`.

Before you begin

If you are using a third party reverse proxy server other than Tivoli WebSEAL, you must provide an implementation class for the product interceptor interface for your proxy server. This article describes the `com.ibm.websphere.security.TrustAssociationInterceptor.java` interface that you must implement.

Note: The Trust Association Interceptor (TAI) interface (`com.ibm.wsspi.security.tai.TrustAssociationInterceptor`) supports several new features and is different from the existing `com.ibm.websphere.security.TrustAssociationInterceptor` interface.

1. Define the interceptor class method. WebSphere Application Server provides the interceptor Java interface, `com.ibm.websphere.security.TrustAssociationInterceptor`, which defines the following methods:

- **public boolean isTargetInterceptor(HttpServletRequest req)** creates `WebTrustAssociationException`;

The `isTargetInterceptor` method determines whether the request originated with the proxy server associated with the interceptor. The implementation code must examine the incoming request object and determine if the proxy server forwarding the request is a valid proxy server for this interceptor. The result of this method determines whether the interceptor processes the request or not.

- **public void validateEstablishedTrust (HttpServletRequest req)** creates `WebTrustAssociationException`;

The `validateEstablishedTrust` method determines if the proxy server from which the request originated is trusted or not. This method is called after the `isTargetInterceptor` method. The implementation code must authenticate the proxy server. The authentication mechanism is proxy-server specific. For example, in the product implementation for the WebSEAL server, this method retrieves the basic authentication information from the HTTP header and validates the information against the user registry used by WebSphere Application Server. If the credentials are invalid, the code creates the `WebTrustAssociationException`, indicating that the proxy server is not trusted and the request is to be denied.

- **public String getAuthenticatedUsername(HttpServletRequest req)** creates `WebTrustAssociationException`;

The `getAuthenticatedUsername` method is called after trust is established between the proxy server and WebSphere Application Server. The product has accepted the proxy server authentication of the request and must now authorize the request. To authorize the request, the name of the original requestor must be subjected to an authorization policy to determine if the requestor has the necessary privilege. The implementation code for this method must extract the user name from the HTTP request header and determine if that user is entitled to the requested resource. For example, in the product implementation for the WebSEAL server, the method looks for an `iv-user` attribute in the HTTP request header and extracts the user ID associated with it for authorization.

2. Configuring the interceptor. To make an interceptor configurable, the interceptor must extend `com.ibm.websphere.security.WebSphereBaseTrustAssociationInterceptor`. Implement the following methods:

public int init (java.util.Properties props);

The `init(Properties)` method accepts a `java.util.Properties` object, which contains the set of properties required to initialize the interceptor. All the properties set for an interceptor (by using the **Custom Properties** link for that interceptor or using scripting) is sent to this method. The interceptor then can use these properties to initialize itself. For example, in the product implementation for the WebSEAL server, this method reads the hosts and ports so that a request coming in can be verified to originate from trusted hosts and ports. A return value of `0` implies that the interceptor initialization is successful. Any other value implies that the initialization is not successful and the interceptor is ignored.

Applicability of the following list

If a previous implementation of the trust association interceptor returns a different error status you can either change your implementation to match the expectations or make one of the following changes:

- Add the `com.ibm.websphere.security.trustassociation.initStatus` property in the trust association interceptor custom properties. Set the property to the value that indicates that the interceptor is successfully initialized. All of the other possible values imply failure. In case of failure, the corresponding trust association interceptor is not used.
- Add the `com.ibm.websphere.security.trustassociation.ignoreInitStatus` property in the trust association interceptor custom properties. Set the value of this property to **true**, which tells WebSphere Application Server to ignore the status of this method. If you add this property to the custom properties, WebSphere Application Server does not check the return status, which is similar to previous versions of WebSphere Application Server.

public void cleanup ();

This method is called when the application server is stopped. It is used to prepare the interceptor for termination.

public void setVersion (String s);

This method is optional. The method is used to set the version and is for informational purpose only. The default value is Unspecified.

You must configure the following methods implemented by the custom interceptor implementation. **This listing only shows the methods and does not include any implementation.**

```
*****
import java.util.*;
import javax.servlet.http.HttpServletRequest;
import com.ibm.websphere.security.*;

public class myTAImpl extends WebSphereBaseTrustAssociationInterceptor
    implements TrustAssociationInterceptor
{
    public myTAImpl ()
    {
    }

    public boolean isTargetInterceptor (HttpServletRequest req)
        creates WebTrustAssociationException
    {
        //return true if this is the target interceptor, else return false.
    }

    public void validateEstablishedTrust (HttpServletRequest req)
        creates WebTrustAssociationFailedException
    {
        //validate if the request is from the trusted proxy server.
        //throw exception if the request is not from the trusted server.
    }

    public String getAuthenticatedUsername (HttpServletRequest req)
        creates WebTrustAssociationUserException
    {
        //Get the user name from the request and if the user is
        //entitled to the requested resource
        //return the user. Otherwise, throw the exception
    }

    public int init (Properties props)
    {
        //initialize the implementation. If successful return 0, else return -1.
    }

    public void cleanup ()
    {
        //Cleanup code.
    }
}
*****
```

Note: If the `init(Properties)` method is implemented as described previously in your custom interceptor, this note does not apply to your implementation, and you can move on to the next step. Previous versions of `com.ibm.websphere.security.WebSphereBaseTrustAssociationInterceptor` include the `public int init (String propsfile)` method. This method is no longer required since the

interceptor properties are not read from a file. The properties are now entered in the administrative console **Custom Properties** link of the interceptor using the administrative console or scripts. These properties then are made available to your implementation in the `init(Properties)` method. However, for backward compatibility, the `init(String)` method still is supported. The `init(String)` method is called by the default implementation of `init(Properties)` as shown in the following example.

```
// Default implementation of init(Properties props) method. A Custom
// implementation should override this.
public int init (java.util.Properties props)
{
    String type =
        props.getProperty("com.ibm.websphere.security.trustassociation.types");
    String classfile=
        props.getProperty("com.ibm.websphere.security.trustassociation."
        +type+".config");
    if (classfile != null && classfile.length() > 0 ) {
        return init(classfile);
    } else {
        return -1;
    }
}
```

Change your implementation to implement the `init(Properties)` method instead of relying on `init(String propsfile)` method. As shown in the previous example, this default implementation reads the properties to load the property file. The `com.ibm.websphere.security.trustassociation.types` property gets the file containing the properties by concatenating `.config` to its value.

Note: The `init(String)` method still works if you want to use it instead of implementing the `init(Properties)` method. The only requirement is that the file name containing the custom trust association properties should now be entered using the **Custom Properties** link of the interceptor in the administrative console or by using scripts. You can enter the property using *either* of the following methods. The first method is used for backward compatibility with previous versions of WebSphere Application Server.

Method 1:

The same property names used in the previous release are used to obtain the file name. The file name is obtained by concatenating the `.config` to the `com.ibm.websphere.security.trustassociation.types` property value.

If the file name is called `myTAI.properties` and is located in the `app_server_root/properties` directory, set the following properties:

- `com.ibm.websphere.security.trustassociation.types = myTAItype`
- `com.ibm.websphere.security.trustassociation.myTAItype.config = app_server_root/properties/myTAI.properties`

Method 2:

You can set the `com.ibm.websphere.security.trustassociation.initPropsFile` property in the trust association custom properties to the location of the file. For example, set the following property:

- `com.ibm.websphere.security.trustassociation.initPropsFile= app_server_root/properties/myTAI.properties`

Type the previous code as one continuous line.

The location of the properties file is fully qualified (for example, `app_server_root/properties/myTAI.properties`). Because the location can be different in a Network Deployment environment, use variables such as `${USER_INSTALL_ROOT}` to refer to the WebSphere Application Server installation directory. For example, if the file name is called `myTAI.properties` and it is located in the `app_server_root/properties` directory, then set the following properties:

3. Compile the implementation once you have implemented it. For example, `app_server_root/java/bin/javac -classpath install_root/plugins/com.ibm.ws.runtime.jar;<install_root>/lib/j2ee.jar myTAIImpl.java`

- a. Copy the class file to a location in the class path (preferably the `app_server_root/lib/ext` directory).

This file needs to be copied to each node class path and each cell class path.

The classes need to be copied to the `app_server_root/lib/ext` directory because the directory is referenced at the beginning of the startup. For the most part, compilation is dependant on this directory being loaded first using the bootstrap.

- b. Restart all the servers.
4. Delete the default WebSEAL interceptor in the administrative console and click **New** to add your custom interceptor. Verify that the class name is dot separated and appears in the class path.
5. Click the **Custom Properties** link to add additional properties that are required to initialize the custom interceptor. These properties are passed to the `init(Properties)` method of your implementation when it extends the `com.ibm.websphere.security.WebSphereBaseTrustAssociationInterceptor` as described in the previous step.
6. Save and synchronize (if applicable) the configuration.
7. Restart the servers for the custom interceptor to take effect.

Example

Refer to the Security: Resources for Learning article for a reference to an example of a custom interceptor.

Trust association interceptor support for Subject creation

The trust association interceptor (TAI) `com.ibm.wsspi.security.tai.TrustAssociationInterceptor` interface supports several features that are different from the existing `com.ibm.websphere.security.TrustAssociationInterceptor` interface.

The TAI interface supports a multiphase, negotiated authentication process. For example, some systems require a challenge response protocol back to the client. The two key methods in this interface are:

Key method name

```
public boolean isTargetInterceptor (HttpServletRequest req)
```

The `isTargetInterceptor` method determines whether the request originated with the proxy server that is associated with the interceptor. The implementation code must examine the incoming request object and determine if the proxy server that forwards the request is a valid proxy server for this interceptor. The result of this method determines whether the interceptor processes the request.

Method result

A `true` value tells WebSphere Application Server to have the TAI handle the request.

A `false` value, tells WebSphere Application Server to ignore the TAI.

Key method name

```
public TAIResult negotiateValidateandEstablishTrust (HttpServletRequest req, HttpServletResponse res)
```

The `negotiateValidateandEstablishTrust` method determines whether to trust the proxy server from which the request originated. The implementation code must authenticate the proxy server. The authentication mechanism is proxy-server specific. For example, in the product implementation for the WebSEAL server, this method retrieves the basic authentication information from the HTTP header and validates the information against the user registry that WebSphere Application Server uses. If the credentials are not valid, the code creates the `WebTrustAssociationException` exception, which indicates that the proxy server is not trusted and the request is denied. If the credentials are valid, the code returns a `TAIResult` result, which indicates the status of the request processing with the client identity (Subject and principal name) to use for authorizing the Web resource.

Method result

Returns a TAIResult result, which indicates the status of the request processing. You can query the Request object and modify the Response object can be modified.

The TAIResult class has three static methods for creating a TAIResult result. The TAIResult create methods take an int type as the first parameter. WebSphere Application Server expects the result to be a valid HTTP request return code and is interpreted in one of the following ways:

- If the value is HttpServletResponse.SC_OK, this response tells WebSphere Application Server that the TAI completed its negotiation. The response also tells WebSphere Application Server to use the information in the TAIResult result to create a user identity.
- Other values tell WebSphere Application Server to return the TAI output, which is placed into the HttpServletResponse response, to the Web client. Typically, the Web client provides additional information and then places another call to the TAI.

The created TAIResults results have the following meanings:

TAIResult	Explanation
public static TAIResult create(int status);	Indicates a status to WebSphere Application Server. The status cannot be SC_OK because the identity information is provided.
public static TAIResult create(int status, String principal);	Indicates a status to WebSphere Application Server and provides the user ID or the unique ID for this user. WebSphere Application Server creates credentials by querying the user registry.
public static TAIResult create(int status, String principal, Subject subject);	Indicates a status to WebSphere Application Server, the user ID or the unique ID for the user, and a custom Subject. If the Subject contains a hashtable, the principal is ignored. The contents of the Subject become part of the eventual user Subject.

All of the following examples are within the negotiateValidateandEstablishTrust method of a TAI.

The following code sample indicates that additional negotiation is required:

```
// Modify the HttpServletResponse object
// The response code is meaningful only on the client
return TAIResult.create(HttpServletResponse.SC_CONTINUE);
```

The following code sample indicates that the TAI determined the user identity. WebSphere Application Server receives the user ID only and queries the user registry for additional information:

```
// modify the HttpServletResponse object
return TAIResult.create(HttpServletResponse.SC_OK, userid);
```

The following code sample indicates that the TAI determined the user identity. WebSphere Application Server receives the complete user information that is contained in the hashtable. For more information on the hashtable, see “Configuring inbound identity mapping” on page 435. In this code sample, the hashtable is placed in the public credential portion of the Subject:

```
// create Subject and place Hashtable in it
Subject subject = new Subject();
subject.getPublicCredentials().add(hashtable);
//the response code is meaningful only the client
return TAIResult.create(HttpServletResponse.SC_OK, "ignored", subject);
```

The following code sample indicates that an authentication failure occurred. WebSphere Application Server fails the authentication request:

```
//log error message
// ....
throw new WebTrustAssociationFailedException("TAI failed for this reason");
```

The following methods are additional methods on the `TrustAssociationInterceptor` interface. These methods are used for initialization, for shutdown, and for identifying the TAI to WebSphere Application Server. For more information, see the Java documentation.

Method name

```
public int initialize(Properties props)
```

Method result

This method is called during TAI initialization and is called only if custom properties are configured for the interceptor.

Method name

```
public String getVersion()
```

Method result

This method returns the version of the TAI.

Method name

```
public String getType()
```

Method result

This method returns the type of the TAI.

Method name

```
public void cleanup()
```

Method result

This method is called when stopping the WebSphere Application Server process. Stopping the WebSphere Application Server process provides an opportunity for the TAI to perform any necessary cleanup. This method is not necessary if cleanup is not required.

Enabling a plugpoint for custom password encryption

Two properties govern the protection of passwords. By configuring these two properties, you can enable a plugpoint for custom password encryption.

Before you begin

To view an example code sample that illustrates the `com.ibm.wsspi.security.crypto.CustomPasswordEncryption` interface, see “Plug point for custom password encryption” on page 835.

About this task

The encryption method is called for password processing whenever the custom class is configured and custom encryption is enabled. The decryption method is called whenever the custom class is configured and the password contains the `{custom:alias}` tag. The `custom:alias` tag is stripped prior to decryption.

1. To enable custom password encryption, you must configure two properties:
 - **com.ibm.wsspi.security.crypto.customPasswordEncryptionClass** - Defines the custom class that implements the `com.ibm.wsspi.security.crypto.CustomPasswordEncryption` password encryption interface.
 - **com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled** - Defines when the custom class is used for default password processing. When the `passwordEncryptionEnabled` option is not specified or set to `false`, and the `passwordEncryptionClass` class is specified, the decryption method is called whenever a `{custom:alias}` tag still exists in the configuration repository.

2. To configure custom password encryption, configure both of these properties in the `security.xml` file. The custom encryption class (`com.acme.myPasswordEncryptionClass`) must be placed in a Java archive (JAR) file in the `#{WAS_INSTALL_ROOT}/classes` directory in all WebSphere Application Server processes. Every configuration document that contains a password (`security.xml` and any application bindings that contain RunAs passwords), must be saved before all of the passwords become encrypted with the custom encryption class.
3. If the custom implementation class defaults to the `com.ibm.wsspi.security.crypto.CustomPasswordEncryptionImpl` interface, and this class is present in the class path, then encryption is enabled by default. This simplifies the enablement process for all nodes. It is not necessary to define any other properties except for those that the custom implementation requires. To disable encryption, but still use this class for decryption, specify the following class.
 - `com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=false`

What to do next

Whenever a custom encryption class encryption operation is called, and it creates a run-time exception or a defined `PasswordEncryptException` exception, the WebSphere Application Server runtime uses the `{xor}` algorithm to encode the password. This encoding prevents the storage of the password in plain text. After the problem with the custom class has been resolved, it automatically encrypts the password the next time the configuration document is saved.

When a RunAs role is assigned a user ID and password, it currently is encoded using the WebSphere Application Server encoding function. Therefore, after the custom plug point is configured to encrypt the passwords, it encrypts the passwords for the RunAs bindings as well. If the deployed application is moved to a cell that does not have the same encryption keys, or the custom encryption is not yet enabled, a login failure results because the password is not readable.

One of the responsibilities of the custom password encryption implementation is to manage the encryption keys. This class must decrypt any password that it encrypted. Any failure to decrypt a password renders that password to be unusable, and the password must be changed in the configuration. All encryption keys must be available for decryption there and no passwords are left using those keys. The master secret must be maintained by the custom password encryption class to protect the encryption keys.

You can manage the master secret by using a stash file for the keystore, or by using a password locator that enables the custom encryption class to locate the password so that it can be locked down.

Plug point for custom password encryption

A plug point for custom password encryption can be created to encrypt and decrypt all passwords in WebSphere Application Server that are currently encoded or decoded using Base64-encoding.

The implementation class of this plug point has the responsibility for managing keys, determining the encryption algorithm to use, and for protecting the master secret. The WebSphere Application Server runtime stores the encrypted passwords in their existing locations, preceded with `{custom:alias}` tags instead of `{xor}` tags. The custom part of the tag indicates that it is a custom algorithm. The alias part of the tag is specified by the custom implementation, which helps to indicate how the password is encrypted. The implementation can include the key alias, encryption algorithm, encryption mode, or encryption padding.

A custom provider of this plug point must implement an interface that is designed to encrypt and decrypt passwords. The interface is called by the WebSphere Application Server runtime whenever the custom plug point is enabled. The custom algorithm becomes one of the supported algorithms when the plug point is enabled. Other supported algorithms include `{xor}` (standard base64 encoding) and `{os400}` which is used on the iSeries platform.

The following example illustrates the `com.ibm.wsspi.security.crypto.CustomPasswordEncryption` interface:

```
package com.ibm.wsspi.security.crypto;
public interface CustomPasswordEncryption
{
    /**
     * The encrypt operation takes a UTF-8 encoded String in the form of a byte[].
     * The byte[] is generated from String.getBytes("UTF-8").
     * An encrypted byte[] is returned from the implementation in the EncryptedInfo
     * object. Additionally, a logical key alias is returned in the EncryptedInfo
     * object which is passed back into the decrypt method to determine which key was
     * used to encrypt this password. The WebSphere Application Server runtime has
     * no knowledge of the algorithm or the key used to encrypt the data.
     *
     * @param byte[]
     * @return com.ibm.wsspi.security.crypto.EncryptedInfo
     * @throws com.ibm.wsspi.security.crypto.PasswordEncryptException
     */
    public EncryptedInfo encrypt (byte[] decrypted_bytes) throws PasswordEncryptException;

    /**
     * The decrypt operation takes the EncryptedInfo object containing a byte[]
     * and the logical key alias and converts it to the decrypted byte[]. The
     * WebSphere Application Server runtime converts the byte[] to a String
     * using new String (byte[], "UTF-8");
     *
     * @param com.ibm.wsspi.security.crypto.EncryptedInfo
     * @return byte[]
     * @throws com.ibm.wsspi.security.crypto.PasswordDecryptException
     */
    public byte[] decrypt (EncryptedInfo info) throws PasswordDecryptException;

    /**
     * The following is reserved for future use and is currently not
     * called by the WebSphere Application Server runtime.
     *
     * @param java.util.HashMap
     */
    public void initialize (java.util.HashMap initialization_data);
}
```

The `com.ibm.wsspi.security.crypto.EncryptedInfo` class contains the encrypted bytes with the user-defined alias that is associated with the encrypted bytes. This information is passed back into the encryption method to help determine how the password was originally encrypted.

```
package com.ibm.wsspi.security.crypto;
public class EncryptedInfo
{
    private byte[] bytes;
    private String alias;

    /**
     * This constructor takes the encrypted bytes and a keyAlias as parameters.
     * This constructor is used to pass to or from the WebSphere Application Server
     * runtime to enable the runtime to associate the bytes with a specific key that
     * is used to encrypt the bytes.
     */

    public EncryptedInfo (byte[] encryptedBytes, String keyAlias)
    {
        bytes = encryptedBytes;
        alias = keyAlias;
    }

    /**
     * This command returns the encrypted bytes.
     */
}
```

```

*
* @return byte[]
*/
public byte[] getEncryptedBytes()
{
    return bytes;
}

/**
 * This command returns the key alias. The key alias is a logical string that is
 * associated with the encrypted password in the model. The format is
 * {custom:keyAlias}encrypted_password. Typically, just the key alias is placed
 * here, but algorithm information can also be returned.
 *
 * @return String
 */
public String getKeyAlias()
{
    return alias;
}
}

```

The encryption method is called for password processing whenever the custom class is configured and custom encryption is enabled. The decryption method is called whenever the custom class is configured and the password contains the {custom:alias} tag . The custom:alias tag is stripped prior to decryption. For more information, see Enabling custom password encryption.

Chapter 9. Auditing the security infrastructure

You can use the Auditing Facility to track and archive auditable events to ensure the integrity of your system.

Before you begin

Before enabling the security auditing subsystem, you must enable global security in your environment.

About this task

Note: The security auditing subsystem has been introduced as part of the security infrastructure. The primary responsibility of the security infrastructure is to prevent unauthorized access and usage of resources. Utilizing security auditing has two primary goals:

- Confirming the effectiveness and integrity of the existing security configuration.
- Identifying areas where improvement to the security configuration might be needed.

Security auditing achieves these goals by providing the infrastructure that allows you to implement your code to capture and store supported auditable security events. During run time, all code other than the Java EE 5 application code is considered to be trusted. Each time a Java EE 5 application accesses a secured resource, any internal application server process with an audit point included can be recorded as an auditable event.

The security auditing subsystem has the ability to capture the following types of auditable events:

- Authentication
- Authorization
- Principal/Credential Mapping
- Audit policy management
- Administrative configuration management
- User registry and identity management
- Delegation

Note: Audit instrumentation has not been included in the Web services client run time.

These types of events can be recorded into audit log files. Each audit log has the option to be signed and encrypted to ensure data integrity. These audit log files can be analyzed to discover breaches over the existing security mechanisms and to discover potential weaknesses in the current security infrastructure. Security event audit records are also useful for providing evidence of accountability and nonrepudiation as well as vulnerability analysis. The security auditing configuration provides four default filters, a default audit service provider, and a default event factory. The default implementation write to a binary text-file based log. Use this topic to customize your security auditing subsystem.

1. “Enabling the security auditing subsystem” on page 840

Security auditing will not be performed unless the audit security subsystem has been enabled. Global security must be enabled for the security audit subsystem to function, as no security auditing occurs if global security is not also enabled.

2. “Administrative roles” on page 518

A user with the auditor role is required to enable and configure the security auditing subsystem. It is important to require strict access control for security policy management. The auditor role has been created providing granularity to allow for separation of the auditing role from the authority of the administrator. When Security Auditing is initially enabled, the cell administrator has auditor privileges. If the environment requires separation of privileges, then changes will need to be made to the default role assignments.

3. “Creating security auditing event type filters” on page 846
You can configure event type filters to only record a specific subset of auditable event types in your audit logs. Filtering the event types that are recorded makes for easier analysis of your audit records by ensuring only those records important to your environment are archived.
4. Configuring the audit service provider.
The audit service provider is used to format the audit data object that was passed to it before outputting the data to a repository. A default audit service provider implementation is included. See “Configuring the default audit service providers for security auditing” on page 856 for more details on the default implementation. A third party implementation can also be coded and used. See “Configuring a third party audit service providers for security auditing” on page 858 for more details on this implementation.
5. “Configuring audit event factories for security auditing” on page 860
The audit event factory gathers the data associated with the auditable events and creates an audit data object. The audit data object is then sent to the audit service provider to be formatted and recorded to the repository.
6. “Protecting your security audit data” on page 863
It is important to secure and ensure the data integrity of the recorded audit data. To ensure that access to the data is restricted and tamper proof, you can encrypt and sign your audit data.
7. “Configuring security audit subsystem failure notifications” on page 853
Notifications can be enabled to generate alerts when the security auditing subsystem experiences a failure. Notifications can be configured to record an alert in the System logs or can be configured to send an alert via e-mail to a specified list of recipients.

Results

After successfully completing this task, you audit data will be recorded for the selected auditable events that were specified in the configuration.

What to do next

After configuring security auditing, you can analyze your audit data for potential weaknesses in the current security infrastructure and to discover security breaches that may have occurred over the existing security mechanisms. You can also use the security auditing subsystem to provide data for problem determination. If the default audit service provider was selected, the resulting binary audit log file can be read using the Audit Reader.

Related tasks

“Configuring security auditing using scripting” on page 1060
Security auditing provides tracking and archiving of auditable events. This topic uses the wsadmin tool to enable and administer your security auditing configurations.

Enabling the security auditing subsystem

Security auditing will not be performed unless the audit security subsystem has been enabled. Global security must be enabled for the security audit subsystem to function, as no security auditing occurs if global security is not also enabled.

Before you begin

Before enabling security auditing subsystem, enable global security in your environment.

About this task

The recording of auditable security events is achieved by enabling the security auditing subsystem. Follow these steps to enable the security auditing subsystem.

1. Click **Security > Security auditing**.
2. Select Enable security auditing. The Enable security auditing check box is not selected by default. This check box must be selected to allow security auditing to be performed with the configurations that have been specified in the audit.xml file.

Note: The audit.xml file is used to store the audit subsystem configurations. Changes to the security auditing subsystem should be made with the user interface or the wsadmin utility. This file should not be edited manually.

3. Select the action from the Audit subsystem failure action dropdown menu to be performed when an audit subsystem failure occurs. Notifications configured to warn of a security auditing subsystem failure will not be posted if the No Warning option is selected for this field. If you select either the Log warning or the Terminate server option, then you must also configure a notification for the action to be performed.
4. Select the Auditor ID from the dropdown menu. The auditor role is needed to make changes to the security auditing configurations. By default, when auditing is first enabled, the primary administrator is also given the auditor role. The primary administrator can then add the auditor role to other users. After the auditor role is added to other users, the auditor role can be removed from the administrator to create a separation of authority between the auditor and the administrator. The Auditor ID is the user considered to be the primary auditor.
5. Optional: Select Enable verbose auditing. When an auditable event is recorded, a default set of audit data is included in the audit data object and recorded to the repository. An additional set of audit data is made available by enabling verbose auditing.
6. Click **Apply**.
7. Restart the application server. The application server must be restarted before the changes go into effect.

Results

The successful completion of these steps results in the security auditing subsystem being enabled.

What to do next

After enabling the security auditing subsystem, refinements can be made to the configuration. You might want to modify the access control of the audit subsystem to separate the authority of the administrator from the authority of the auditor. If no changes to your access control are needed, then you can configure the types of auditable security events that should be recorded. To configure the types of events that are recorded, click **Event type filters**.

Related tasks

“Enabling security auditing using scripting” on page 1065

Use this task to enable and configure security auditing in your environment with the wsadmin tool. Security auditing provides tracking and archiving of auditable events.

“Configuring security audit subsystem failure notifications” on page 853

Notifications can be generated by a failure of the security audit subsystem. The security audit subsystem notifications can alert auditors that the security audit system is no longer recording auditable security events. Notifications are generated by a failure of the auditing subsystem, they are not related to any auditable security events or event outcome that has occurred. Notifications triggered by an event or an event outcome are not supported.

Security Auditing detail

The Security auditing subsystem can be enabled and configured from this panel, by users assigned the auditor role.

To view this administrative console page, click **Security > Security Auditing**. If Enable security auditing is not selected, then all of the other fields on this panel will be disabled.

Enable security auditing

The Enable security auditing check box allows users to enable or disable Security Auditing. By default, Security Auditing will not be enabled. This field corresponds with the auditEnabled field in the audit.xml file.

Audit subsystem failure action

The Audit subsystem failure action setting describes the behavior of the application server in the event of a failure in the auditing subsystem. Audit Notifications must be configured in order for notifications of a failure in the audit subsystem to be logged. If security auditing is not enabled, then these actions will not be performed. Failures can include an error in the interface or in the event processing. By default, the audit subsystem failure action setting is set to No warning.

The Audit subsystem failure action dropdown menu has the following options:

- No warning
The No warning action specifies that the auditor will not be notified of a failure in the audit subsystem. The product will continue processing but audit reporting will be disabled.
- Log warning
The Log warning action specifies that the auditor will be notified of a failure in the audit subsystem. The product will continue processing but audit reporting will be disabled.
- Terminate server
The Terminate server action specifies the application server to gracefully quiesce when an unrecoverable error occurs in the auditing subsystem. If e-mail notifications are configured, the auditor will be sent a notification that an error has occurred. If logging to the system log is configured, the notification of the failure will be logged to the system file.

Primary auditor user name

The Primary auditor user name dropdown menu defines a valid user which exists in the current user registry and for whom the auditor role has been given. By default, this field is blank and is a required field.

Enable verbose auditing

The Enable verbose auditing option determines the amount of audit data that is reported in an audit record. Verbose mode captures all the auditable data points, whereas not enabling verbose mode captures only a subset of the available data. This option is disabled by default.

Context object fields

Each auditable event has an associated set of information that is available for logging. This information is grouped into specific context objects. The context objects that are available for logging a specific event are specified by the event type. This topic details the information that exists for each context object and specifies whether the information is logged by default or is only logged when the verbose logging option is enabled.

The SessionContextObj object

Table 24. SessionContextObj fields

Field	Type	Description	Default or Verbose logging
sessionId	String	An identifier for the user session	Default
remoteAddr	String	The IP address for the remote host	Default
remotePort	String	The port of the remote host	Default
remoteHost	String	The host name of the remote host	Default

The PropagationContextObj object

Table 25. PropagationContextObj fields

Field	Type	Description	Default or Verbose logging
firstCaller	String	The identity of the first user in the caller list	Default
callerList	String array	A list of names representing the identities of the users	Verbose

The RegistryContextObj object

Table 26. RegistryContextObj fields

Field	Type	Description	Default or Verbose logging
type	String	The type of user registry being used, such as LDAP or AIX	Default

The ProcessContextObj object

Table 27. ProcessContextObj fields

Field	Type	Description	Default or Verbose logging
domain	String	The domain to which the user belongs	Verbose
realm	String	The registry partition to which the user belongs	Default

The EventContextObj object

Table 28. EventContextObj fields

Field	Type	Description	Default or Verbose logging
lastEventTrailId	String	The last ID associated with a given transaction	Verbose
eventTrailId	String array	An array of IDs that allow events that belong to a given transaction to be correlated	Default
creationTime	Date	The date an event was created	Default
globalInstanceId	Long	The unique identifier of this event	Default

The DelegationContextObj object

Table 29. DelegationContextObj fields

Field	Type	Description	Default or Verbose logging
delegationType	String	no delegation, simple delegation, method delegation or switch user delegation	Default
roleName	String	The Run as role being used: runAsClient, runAsSpecified, runAsSystem, own ID	Default
identityName	String	Information about the mapped user	Default

The AuthnContextObj object

Table 30. AuthnContextObj fields

Field	Type	Description	Default or Verbose logging
authnType	String	The type of authentication used	Default

The ProviderContextObj object

Table 31. ProviderContextObj fields

Field	Type	Description	Default or Verbose logging
provider	String	The provider of the authentication or authorization service	Default
providerStatus	String	Status of whether the authentication or authorization event processed successfully by the provider	Default

The AuthnMappingContextObj object

Table 32. AuthnMappingContextObj fields

Field	Type	Description	Default or Verbose logging
mappedSecurityDomain	String	The security domain after mapping has occurred	Default
mappedRealm	String	The realm after mapping has occurred	Default
mappedUserName	String	The user name after mapping has occurred	Default

The AuthnTermContextObj object

Table 33. AuthnTermContextObj fields

Field	Type	Description	Default or Verbose logging
terminateReason	String	The reason authentication ended	Default

The AccessContextObj object

Table 34. AccessContextObj fields

Field	Type	Description	Default or Verbose logging
progName	String	The name of the program that was involved in the event	Default
action	String	The action being performed.	Default
registryUserName	String	The name of the user in the registry	Default
appUserName	String	The name of the user within an application	Default
accessDecision	String	The decision of the authorization call	Default

Table 34. AccessContextObj fields (continued)

Field	Type	Description	Default or Verbose logging
resourceName	String	The name of the resource in the context of the application	Default
resourceType	String	The type of resource	Default
resourceUniqueld	Long	The unique identifier of the resource	Default
permissionsChecked	String array	The permissions that were checked during the authorization call	Default
permissionsGranted	String array	The permissions that were granted during the authorization call	Default
rolesChecked	String array	The roles that were checked during the authorization call	Default
rolesGranted	String array	The roles that were granted during the authorization call	Default

The PolicyContextObj object

Table 35. PolicyContextObj fields

Field	Type	Description	Default or Verbose logging
policyName	String	The name of the policy	Default
policyType	String	The type of policy	Default

The KeyContextObj object

Table 36. KeyContextObj fields

Field	Type	Description	Default or Verbose logging
keyLabel	String	The key or certificate label	Default
keyLocation	String	The physical location of the key database	Default
certLifetime	Date	The date when a certificate expires	Default

The CipherContextObj object

Table 37. CipherContextObj fields

Field	Type	Description	Default or Verbose logging
cipherData	Byte array	The cipher data that is captured	Verbose

The MgmtContextObj object

Table 38. MgmtContextObj fields

Field	Type	Description	Default or Verbose logging
mgmtType	String	The type of management operation	Default
mgmtCommand	String	The application-specific command that was performed	Default
targetInfoAttributes	Target Attribute array	Information about one or more secondary objects involved in this operation	Verbose

The ResponseContextObj object

Table 39. ResponseContextObj fields

Field	Type	Description	Default or Verbose logging
url	String	The URL of the HTTP request	Default
httpRequestHeaders	Attributes array	The HTTP request headers provided by the client	Verbose

Table 39. ResponseContextObj fields (continued)

Field	Type	Description	Default or Verbose logging
httpResponseHeaders	Attributes array	The HTTP response headers returned by the server	Verbose

The CustomPropertyContextObj object

Table 40. CustomPropertyContextObj fields

Field	Type	Description	Default or Verbose logging
key	String	The label representing the custom property key name	Verbose
value	Object	The object value of the custom property	Verbose

Creating security auditing event type filters

Event type filters are used to specify the types of auditable security events that are audited. Default event type filters are included with the product, but you can also configure new event type filters to specify a subset of auditable event types to be recorded by the security auditing subsystem.

Before you begin

Before configuring security auditing filters and the rest of the security auditing subsystem, enable global security in your environment. You must be assigned the auditor role to complete this task. Event type filters are used to specify what events are audited. The amount of data that is recorded for each event is specified with the **Enable verbose auditing** check box on the same panel used to enable the auditing subsystem. Navigate to **Security** → **Security auditing** to enable security auditing and determine the data recorded for each event.

About this task

The application server provides the following commonly used event type filters by default in the `audit.xml` template file:

Name	Event name	Outcome of event
DefaultAuditSpecification_1	SECURITY_AUTHN	SUCCESS
DefaultAuditSpecification_2	SECURITY_AUTHN	DENIED
DefaultAuditSpecification_3	SECURITY_RESOURCE_ACCESS	SUCCESS
DefaultAuditSpecification_4	SECURITY_AUTHN	REDIRECT

New event type filters can be created, or the existing default filters can be extended, to capture more event types and outcomes. Use this task to create new event type filters.

1. Click **Security** > **Security Auditing** > **Event type filters** > **New**.
2. Enter the unique name that should be associated with this event type filter configuration in the Name field.
3. Specify the events that should be recorded when this filter is applied:
 - a. Select the events that you want to be audited from the Selectable events list.
 - b. Click **Add >>** to add the selected events to the Enabled events list.
 - c. Select the outcomes that you want to be audited from the Selectable event outcomes list.
 - d. Click **Add >>** to add the selected outcomes to the Enabled event outcomes lists.
4. Click **OK**.

Results

The successful completion of this task results in the creation of an event type filter than can be selected by the audit service providers and audit event factories to gather and record a specific set of auditable security events.

What to do next

After creating an event type filter, the filter must be specified in the audit service provider and the audit event factory to be used to gather or report audit data. The next step in configuring the security auditing subsystem is you should configure an audit service provider to define where the audit data will be archived.

Auditable security events

Auditable security events are security events that have audit instrumentation added to the security run time code to enable them to be recorded. Event filters are configured to specify which auditable security events are recorded to the audit log files.

The following list describes each valid auditable event that you can specify as an enabled event type when creating an event filter:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where there are two user identities involved.
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type will also be used when switching user identities within a given session.

For each audit event type, you must specify an outcome. Valid outcomes include SUCCESS, FAILURE, REDIRECT, ERROR, DENIED, WARNING, and INFO. Not all outcomes are applicable with all event types.

Event type filter settings

The Event type filter settings panel is used by an auditor to manage and create event type filters. Default event type filters have been included, this panel allows additional event type filters to be added. Existing event type filters are also managed using this panel.

To view this administrative console page, click one of the following paths:

- **Security > Security Auditing > Event type filters > *event_type_filter_name*.**
- **Security > Security Auditing > Event type filters > New .**

Name

The Name field specifies the unique name of the event type filter.

Enabled

The state of enablement of the filter is defined by the Enable check box. This field is represented as a boolean value. A value of true specifies that the enable field associated with the audit specification in the audit.xml is set to true. It does not imply that all configured event factories and service providers will be using this filter.

Filters still need to be configured for each event factory and service provider. Filters are enabled by default during configuration. However, if a filter has the enabled checkbox set to false, the filter will not gather or report data for the events and outcomes defined in that filter.

Events to associate with an audit filter

The Events to associate with an audit filter field specifies the auditable security events to be associated with this filter.

- Selectable events:

The Selectable events list displays the available auditable security events. To enable an event for this filter, select the event from the Selectable event outcomes list and then click **Add**.

- Enabled events:

The Enabled events list displays the audit security events that are currently enabled for this filter. To disable an event for this filter, select the event from the Enabled events list and then click **Remove**.

Event outcomes to associate with an audit filter

The Event outcomes to associate with an audit filter field specifies the auditable security event outcomes to be associated with this filter.

- Selectable event outcomes:

The Selectable event outcomes list displays the available auditable security event outcomes. To enable an event outcome for this filter, select the event outcome from the Selectable event outcomes list and then click **Add**.

- Enabled event outcomes:

The Enabled event outcomes list displays the audit security event outcomes that are currently enabled for this filter. To disable an event outcome for this filter, select the event outcome from the Enabled event outcomes list and then click **Remove**.

Event type filters collection

The Event type filters panel displays a listing of all configured audit specifications with their unique names, the state of their enablement, and the event types and event outcomes that are specified for each configuration.

To view this administrative console page, click **Security > Security Auditing > Event type filters**.

Name

The Name field displays the unique name of the event type filter that is being represented.

Enable

The Enable check box species the state of enablement for the filter. This field is represented as a boolean value. A value of true specifies that the enable field associated with the audit specification in the audit.xml is set to true. It does not imply that all configured event factories and service providers will be using this filter. Filters are enabled by default when they are created. Even though it is enabled by default when it is created, the event type filter must be specified for the event factory and the audit service provider before it is actually used,

Filters still need to be configured for each event factory and service provider. A filter that is configured for an event factory or a service provider that has Enabled set to false, will not gather or report data for the events and outcomes defined in that filter.

Events and outcomes

The event types and the event outcomes that are specified by this filter. The specifications are listed in the form event_type:event_outcome and separated by commas if multiple combinations are specified by the event type filter.

Example: Generic Event Interface

This interface is used for processing generic audit events. Other interfaces can be defined which extend this interface to process specific audit event groupings, such as security events, transaction events, or other custom groupings. For WebSphere Application Server version 7.0, only security types of events are supported.

Generic Event Interface

Specific implementations might be developed to handle the data in a particular internal format. When the `buildEvent()` method is called, the implementation must then build the specified base event type using the internal information it has stored. After the information has been stored into a `GenericEvent` instance, the `GenericEvent` interface provides a generic way of handling the event.

```
public interface GenericEvent {

    /** * Property name used to specify the base event type to the
    * {@link GenericEvent#buildEvent} method.
    */
    public static final String BASE_EVENT_TYPE = GenericEvent.class.getName() + ".baseEventType";

    /**
    * Returns the eventType of the event. The eventType distinguishes between these
    * related events.
    * The eventType depends on the particular implementation
    * of the GenericEvent. For example, the Security Event implementation has
    * eventTypes such as SECURITY_AUTHN and SECURITY_AUTHZ.
    * @return eventType - the eventType of the event
    */

    public String getEventType();
    /**
    * Returns the creationTime, the creation time of the event.
    * @return creationTime - the creation time of the event
    */ public Date getCreationTime(); /** * Returns the version, the version of the event.
    * @return version - the version of the event
    */

    public String getVersion (Properties props) throws GenericEventConfigurationException;
    /**
    * Returns the globalInstanceId, which is a globally unique instance
    * identifier for the event.
    * @return globalInstanceId - a globally unique instance identifier for the event
    */

    public Long getGlobalInstanceId();
    /**
    * Verifies whether the event is valid; which depends on the particular
    * implementation of the GenericEvent. If the event is not valid, an
    * GenericEventValidationException error occurs.
    */

    public void validate() throws GenericEventValidationException;
    /**
    * Returns the internally wrapped base event instance after
    * completing and validating the current instance of the GenericEvent.
    * An GenericEvent implementation can maintain its information
    * in any undisclosed internal format. The buildEvent()
    * method that specifies that a specific base event type be built
    * using the internal information. This allows GenericEvent implementations
    * to support multiple base event formats. Thus the GenericEvent implementation
    * provides a layer of abstraction higher than the base event type.
    * @param properties The value of the property BASE_EVENT_TYPE
    * defines the type of the base event * @return the internally wrapped base event instance
    * @throws GenericEventConfigurationException if the base event type is invalid
    * or the JAR files to support that event type are not available.
    * @throws GenericEventCompletionException if event completion has failed.
    * @throws GenericEventValidationException if the validation has failed. This is
    * validation as is performed by the validate() method.
    */

    public Object buildEvent(Properties properties)
        throws GenericEventConfigurationException,
            GenericEventValidationException,
            GenericEventCompletionException;
    /**
    * Returns the wrapped base event instance as a string after
    * completing and validating the current instance of the GenericEvent.
    * An GenericEvent implementation can maintain its information
    * in any undisclosed internal format. It is the buildEventString()
    * method that specifies that a specific base event type be built
    * using the internal information. This allows GenericEvent implementations
    * to support multiple base event formats. Thus the GenericEvent implementation
```

```

* provides a layer of abstraction higher than the base event type.
* @param properties The value of the property BASE_EVENT_TYPE
* defines the type of the base event
* @return the wrapped base event instance as a String
* @throws GenericEventConfigurationException if the base event type is invalid
* or the JAR files to support that event type are not available.
* @throws GenericEventCompletionException if event completion has failed.
* @throws GenericEventValidationException if the validation has failed. This is
* validation as is performed by the validate() method.
*/

```

```

public String buildEventString(Properties properties)
    throws GenericEventConfigurationException,
           GenericEventValidationException,
           GenericEventCompletionException;
}

```

Context objects for security auditing

Each event has an associated set of information that is available for logging. This information is grouped into specific context objects. The context objects that are available for logging a specific event are specified by the event type. All event types have the `sessionContextObj`, `eventContextObj`, `accessContextObj`, `propagationContextObj`, `processContextObj` and `registryContextObj` objects. This topic specifies which additional context objects are available for each event type.

Table 41.

Event Type	Additional Context Objects
SECURITY_AUTHN	authnContextObj, providerContextObj
SECURITY_AUTHN_DELEGATION	delegationContextObj
SECURITY_AUTHN_MAPPING	authnMappingContextObj, providerContextObj
SECURITY_AUTHZ	providerContextObj, policyContextObj
SECURITY_ENCRYPTION	keyContextObj
SECURITY_MGMT_AUDIT	mgmtContextObj
SECURITY_RESOURCE_ACCESS	responseContextObj

For more details on the auditable data that is gather for each of these context objects, see the information for context object fields.

Context object fields

Each auditable event has an associated set of information that is available for logging. This information is grouped into specific context objects. The context objects that are available for logging a specific event are specified by the event type. This topic details the information that exists for each context object and specifies whether the information is logged by default or is only logged when the verbose logging option is enabled.

The SessionContextObj object

Table 42. SessionContextObj fields

Field	Type	Description	Default or Verbose logging
sessionId	String	An identifier for the user session	Default
remoteAddr	String	The IP address for the remote host	Default
remotePort	String	The port of the remote host	Default
remoteHost	String	The host name of the remote host	Default

The PropagationContextObj object

Table 43. PropagationContextObj fields

Field	Type	Description	Default or Verbose logging
firstCaller	String	The identity of the first user in the caller list	Default
callerList	String array	A list of names representing the identities of the users	Verbose

The RegistryContextObj object

Table 44. RegistryContextObj fields

Field	Type	Description	Default or Verbose logging
type	String	The type of user registry being used, such as LDAP or AIX	Default

The ProcessContextObj object

Table 45. ProcessContextObj fields

Field	Type	Description	Default or Verbose logging
domain	String	The domain to which the user belongs	Verbose
realm	String	The registry partition to which the user belongs	Default

The EventContextObj object

Table 46. EventContextObj fields

Field	Type	Description	Default or Verbose logging
lastEventTrailId	String	The last ID associated with a given transaction	Verbose
eventTrailId	String array	An array of IDs that allow events that belong to a given transaction to be correlated	Default
creationTime	Date	The date an event was created	Default
globalInstanceId	Long	The unique identifier of this event	Default

The DelegationContextObj object

Table 47. DelegationContextObj fields

Field	Type	Description	Default or Verbose logging
delegationType	String	no delegation, simple delegation, method delegation or switch user delegation	Default
roleName	String	The Run as role being used: runAsClient, runAsSpecified, runAsSystem, own ID	Default
identityName	String	Information about the mapped user	Default

The AuthnContextObj object

Table 48. AuthnContextObj fields

Field	Type	Description	Default or Verbose logging
authnType	String	The type of authentication used	Default

The ProviderContextObj object

Table 49. ProviderContextObj fields

Field	Type	Description	Default or Verbose logging
provider	String	The provider of the authentication or authorization service	Default
providerStatus	String	Status of whether the authentication or authorization event processed successfully by the provider	Default

The AuthnMappingContextObj object

Table 50. AuthnMappingContextObj fields

Field	Type	Description	Default or Verbose logging
mappedSecurityDomain	String	The security domain after mapping has occurred	Default
mappedRealm	String	The realm after mapping has occurred	Default
mappedUserName	String	The user name after mapping has occurred	Default

The AuthnTermContextObj object

Table 51. AuthnTermContextObj fields

Field	Type	Description	Default or Verbose logging
terminateReason	String	The reason authentication ended	Default

The AccessContextObj object

Table 52. AccessContextObj fields

Field	Type	Description	Default or Verbose logging
progName	String	The name of the program that was involved in the event	Default
action	String	The action being performed.	Default
registryUserName	String	The name of the user in the registry	Default
appUserName	String	The name of the user within an application	Default
accessDecision	String	The decision of the authorization call	Default
resourceName	String	The name of the resource in the context of the application	Default
resourceType	String	The type of resource	Default
resourceUniqueld	Long	The unique identifier of the resource	Default
permissionsChecked	String array	The permissions that were checked during the authorization call	Default
permissionsGranted	String array	The permissions that were granted during the authorization call	Default
rolesChecked	String array	The roles that were checked during the authorization call	Default
rolesGranted	String array	The roles that were granted during the authorization call	Default

The PolicyContextObj object

Table 53. PolicyContextObj fields

Field	Type	Description	Default or Verbose logging
policyName	String	The name of the policy	Default
policyType	String	The type of policy	Default

The KeyContextObj object

Table 54. KeyContextObj fields

Field	Type	Description	Default or Verbose logging
keyLabel	String	The key or certificate label	Default
keyLocation	String	The physical location of the key database	Default
certLifetime	Date	The date when a certificate expires	Default

The CipherContextObj object

Table 55. CipherContextObj fields

Field	Type	Description	Default or Verbose logging
cipherData	Byte array	The cipher data that is captured	Verbose

The MgmtContextObj object

Table 56. MgmtContextObj fields

Field	Type	Description	Default or Verbose logging
mgmtType	String	The type of management operation	Default
mgmtCommand	String	The application-specific command that was performed	Default
targetInfoAttributes	Target Attribute array	Information about one or more secondary objects involved in this operation	Verbose

The ResponseContextObj object

Table 57. ResponseContextObj fields

Field	Type	Description	Default or Verbose logging
url	String	The URL of the HTTP request	Default
httpRequestHeaders	Attributes array	The HTTP request headers provided by the client	Verbose
httpResponseHeaders	Attributes array	The HTTP response headers returned by the server	Verbose

The CustomPropertyContextObj object

Table 58. CustomPropertyContextObj fields

Field	Type	Description	Default or Verbose logging
key	String	The label representing the custom property key name	Verbose
value	Object	The object value of the custom property	Verbose

Configuring security audit subsystem failure notifications

Notifications can be generated by a failure of the security audit subsystem. The security audit subsystem notifications can alert auditors that the security audit system is no longer recording auditable security events. Notifications are generated by a failure of the auditing subsystem, they are not related to any auditable security events or event outcome that has occurred. Notifications triggered by an event or an event outcome are not supported.

Before you begin

Before configuring notifications, enable global security and the security audit subsystem in your environment. You must be assigned the auditor role to complete this task.

About this task

If a problem is experienced with the security audit subsystem, then a notification can be generated. This is an alert that security events are no longer being audited. Notification can be written to the system log file or can be sent to a specified group of users as an e-mail. You are able to configure notifications to alert

the auditor of a problem using both of these methods simultaneously. Notifications are only generated when the Audit subsystem failure action field is set to Log warning or Terminate server.

1. Optional: Click **Security** → **Security Auditing**.
2. Optional: Confirm the Audit subsystem failure action field is set to Log warning or Terminate server. If the Audit subsystem failure action field is set to No warning, then notifications will not be generated.
3. Click **Security** → **Security Auditing** → **Audit monitor** .
4. Under Notifications, Click **New**
5. Enter the name that should be associated with this notification configuration in the Notification name field.
6. Select the Message log check box to specify the failure notifications are recorded in the audit log.
7. Select the e-mail sent to notification list check box to specify that failure notification e-mail should be sent to the addresses listed in the notification list.
8. Enter an e-mail address in the E-mail address to add field This step is not needed if e-mail notifications are not going to be sent.
9. Enter the mail server address in the Outgoing mail (STMP) server address. This step is not needed if e-mail notifications are not going to be sent.
10. Click **Add >>** to add the e-mail address and associated mail server to the e-mail notification list.
11. Repeat steps 5 through 7 for each e-mail address you want to specify in the e-mail notification list.
12. Click **OK**.
13. Select the Enable monitoring check box to turn on audit failure notifications.
14. Select the notification configuration to be used from the Monitor notification dropdown menu.
15. Click **OK**.

Results

After completing this task, a notification will be generated if the security auditing subsystem experiences an unrecoverable error resulting in security events no longer being audited.

What to do next

After configuring notifications, you can analyze your audit data for potential weaknesses in the current security infrastructure and to discover possible security breaches that might have occurred.

Audit notifications cannot be removed using the administrative console. To remove an audit notification you first must run the `deleteAuditNotificationMonitorByRef` or the `deleteAuditNotificationMonitorByName` command. After running one of those commands, remove the audit notification by running the `deleteAuditNotification` command.

Audit monitor collection

Use this page to configure audit subsystem failure notifications. The Auditor monitor panel lists the existing notification configurations and is the gateway for creating new notification configurations and for managing the existing notification configurations.

To view this administrative console page, click **Security** → **Security Auditing** → **Audit monitor**.

Related reference

“Security Auditing detail” on page 842

The Security auditing subsystem can be enabled and configured from this panel, by users assigned the auditor role.

“Audit notification settings”

Use this page to create and manage notification configurations that define how auditors are made aware of audit subsystem failures.

Enable monitoring

Specifies whether to enable or disable notifications. If the check box is selected, then monitoring is enabled. If the check box is not selected, then monitoring is disabled. This check box is disabled by default.

Monitor notification

Specifies the notification configuration that will be used for reporting audit subsystem failures.

Notification name

Specifies a string that uniquely identifies a notification configuration.

Message log

Specifies if the configuration will send failure notifications to the message log file. If the value is true, then failure notifications will be sent to the message log file. If the value is false, then failure notifications will be not be sent to the message log file. When creating a notification, this field is in the form of a check box and is not selected by default.

Send E-mail

Specifies whether an e-mail notification is sent to the addresses listed in the List of e-mail addresses column.

List of e-mail addresses

Specifies the e-mail addresses listed as recipients for e-mail notification in the event of an audit subsystem failure. No e-mail addresses are listed by default. E-mail addresses will appear in this column if they are listed in the notification list in the notification, this applies even when the E-mail sent to notification list check box is not selected in the notification.

Audit notification settings

Use this page to create and manage notification configurations that define how auditors are made aware of audit subsystem failures.

To view this administrative console page, click **Security** → **Security Auditing** → **Audit monitor** → **New**.

Notification name

Specifies a string that uniquely identifies a notification configuration.

Message log

Specifies if the configuration will send failure notifications to the message log file. If the check box is selected, then failure notifications will be sent to the message log file. If the check box is not selected, then failure notifications will be not be sent to the message log file. This check box is not selected by default.

Send secure e-mails

E-mail sent to the notification list

Specifies whether the configuration will send a failure notification to the recipients listed in the notification list. If the check box is selected, then failure notifications will be sent to the recipients in the notification list.

If the check box is not selected, then failure notifications will be not be sent to the recipients in the notification list. This check box is not selected by default.

E-mail address to add

Specifies the e-mail address to be added to the notification list to received failure notification e-mails. To add a recipient to the notification list, this field and the Outgoing mail (SMTP) server field must both be completed before you click the **Add**.

Outgoing mail (SMTP) server

Specifies the SMTP server to be used with this e-mail address. If no server is specified, then the e-mail realm will be used.

Configuring the default audit service providers for security auditing

The audit service provider is used to format the audit data object that was sent by the audit event factory. After being formatted, the audit data is recorded to the repository defined in the audit service provider configuration.

Before you begin

Before configuring the audit service provider, enable global security in your environment.

About this task

This task configures the audit service provider used to record generated audit records.

1. Click **Security** → **Security Auditing** → **Audit service provider**.
2. Click **New** and then select **Binary file based emitter**.
3. Enter the unique name that should be associated with this audit service provider in the Name field.
4. Enter the file location of the binary log file in the Audit log file location field.
5. Optional: Enter the maximum size allowed for a single binary log file in the Audit log file size field.
This field is specified in megabytes. After the maximum audit file size is reached, a new audit file will be created or an existing audit file will be overwritten. If the maximum number of audit log files has not been set, the default maximum file value used is 10 megabytes. There is no audit archiving utility included with the product. You are responsible for the archiving of your audit data.
6. Optional: In the Maximum number of audit log files field, enter the maximum number of audit logs to be stored before the oldest is overwritten.
The default value for this field is 100. The value of 100 is also used if the field is empty.
7. Select the filters to be used by this audit service provider. The Selectable filter list consists of a list of the configured filters that have been configured and are currently enabled.
 - a. Select the filters that should be audited from the Selectable filter list.
 - b. Click **Add >>** to add the selected filters to the Enabled filter list.
8. Click **Apply**.

Results

After completing these steps, your audit data will be sent to the specified repository in the format required by that repository.

What to do next

After creating an audit service provider, the audit service provider must be associated with an audit event factory provide the audit data objects to the audit service provider. Next you should configure an audit event factory.

Audit service provider collection

The Audit service provider panel displays a listing of all configured audit service provider implementations. Using this panel, a user can define a new audit service provider implementation, delete an existing implementation, and display or modify the fields associated with an existing implementation.

To view this administrative console page, click **Security > Security Auditing > Audit service provider**.

By default, the audit.xml will contain the IBM audit service provider implementation which emits audit records to a binary file-based text file. This implementation is used for Binary file-based audit service provider configurations. For each existing audit service provider in the list on this panel, the unique name, type and event formatting class associated with the audit service provider will be displayed.

Name

The Name field is the unique name associated with the audit service provider implementation.

Type

The Type field specifies if the implementation is a binary file-based implementation, SMF implementation or a third party implementation.

Event formatting module class name

The event formatting class is a class used to format the generic event data object into a format that is specific to the audit service provider implementation. For example, a third party audit service provider implementation might have an event formatting class that takes the generic event and translates it into XML data. There is no Event formatting module class for binary file-based implementations nor for SMF implementations.

Audit service provider settings

Use this page to define the implementation details of the audit service provider. There are three types of audit service providers: binary file-based, third party and SMF.

To view this administrative console page, click on of the following paths:

- **Security > Security auditing > Audit service provider > *audit_service_provider_name***.
- **Security > Security auditing > Audit service provider > New > Binary File-based emitter**.
- **Security > Security auditing > Audit service provider > New > Third party emitter**.

Name

Specifies the unique name associated with the audit service provider.

Third party emitter class name

Specifies the name of the class for this implementation. This field is only present for Third party emitter implementations.

Audit file location

Specifies the path to the binary log file.

Audit file size

Specifies the maximum size of a single binary log file. This value is defined in megabytes.

Maximum number of audit log files

Specifies the maximum number of binary log file to create before the oldest is replaced.

Event formatting module class name

Specifies a class used to format the generic event into a format that is specific to the audit service provider implementation. For example, a third party audit service provider implementation might have an event formatting class that takes the generic event and translates it into XML data.

Selectable filters

Specifies the available event filters. To enable a filter for an implementation, select the filter from the Selectable event filters list and then click >.

Enabled filters

Specifies the event filters that are currently enabled for an implementation. To disable a filter for an implementation, select the filter from the Enabled filters list and then click <.

Custom properties

Specifies any custom properties that might be used to add properties to a third party implementation. Custom properties are not available for binary file-based implementations or SMF implementations.

- Name
- Value

Example: Base Generic Emitter Interface

The Base Generic Emitter interface defines how audit events are emitted. Other interfaces can exist to extend this interface and to process specific audit events groupings, such as security events, transactional events, or some other custom grouping. Use this interface to create a custom implementation of the emitter.

Base Generic Emitter Interface

```
/**
 * This is the interface for the event emitter. Event sources use this interface
 * to send events to an event service.
 */
public interface BaseGenericEmitter {
    /**
     * Sends an event to the configured GenericEmitter implementation.
     *
     * @param event The event to be sent to the event service.
     * This value cannot be null.
     * @return The global instance ID of the event that was built.
     * @exception GenericEmitterException If an error occurs during emitter processing.
     * @exception IllegalArgumentException If the event parameter is null.
     */
    public String sendEvent(GenericEvent event) throws
        GenericEventException;
    /**
     * Sends an array of events to the configured GenericEmitter implementation.
     * @param events The event array to be sent to the event service.
     * This value cannot be null.
     * @return The global instance IDs of the events that were built.
     * @exception GenericEmitterException If an error occurs during emitter processing.
     * @exception IllegalArgumentException If the events parameter is null.
     */
    public String[] sendEvents(GenericEvent events[]) throws
        GenericEventException;
    /**
     * Causes the emitter to release all resources that are owned by this
     * object and its dependents.
     * Subsequent calls to this method have no effect.
     *
     * @throws GenericEmitterException If the emitter does release the
     * held resources.
     * resources.
     * @throws GenericEventException If any other error occurs when releasing resources.
     */
    public void close() throws
        GenericEventException;
}
```

Configuring a third party audit service providers for security auditing

The audit service provider is used to format the audit data object that was sent by the audit event factory. In addition to the default audit service provider, you may use a third party implementation as your audit service provider.

Before you begin

Before configuring the audit service provider, enable global security in your environment.

About this task

This task configures the audit service provider used to record generated audit records.

1. Click **Security** → **Security Auditing** → **Audit service provider**.
2. Click **New** and then select **Third party emitter**.
3. Enter the unique name that should be associated with this audit service provider in the Name field.
4. Enter the Third party emitter class name.
5. Enter the Event formatting module class name. This field specifies the class used to format the generic event into a format that is specific to the audit service provider implementation. For example, your implementation might have an event formatting class that takes the generic event and translates it into XML data.
6. Select the filters to be used by this audit service provider. The Selectable filter list consists of a list of the configured filters that have been configured and are currently enabled.
 - a. Select the filters that should be audited from the Selectable filter list.
 - b. Click **Add >>** to add the selected filters to the Enabled filter list.
7. Optional: Enter any custom properties that you included in your third party emitter code.
8. Click **Apply**.

Results

After completing these steps, your audit data will be sent to the specified repository in the format required by that repository.

What to do next

After creating an audit service provider, the audit service provider must be associated with an audit event factory provide the audit data objects to the audit service provider. Next you should configure an audit event factory.

Example: Base Generic Emitter Interface

The Base Generic Emitter interface defines how audit events are emitted. Other interfaces can exist to extend this interface and to process specific audit events groupings, such as security events, transactional events, or some other custom grouping. Use this interface to create a custom implementation of the emitter.

Base Generic Emitter Interface

```
/**
 * This is the interface for the event emitter. Event sources use this interface
 * to send events to an event service.
 */
public interface BaseGenericEmitter {
    /**
     * Sends an event to the configured GenericEmitter implementation.
     *
     * @param event The event to be sent to the event service.
     * This value cannot be null.
     * @return The global instance ID of the event that was built.
     * @exception GenericEmitterException If an error occurs during emitter processing.
     * @exception IllegalArgumentException If the event parameter is null.
     */
    public String sendEvent(GenericEvent event) throws
        GenericEventException;
    /**
     * Sends an array of events to the configured GenericEmitter implementation.
     * @param events The event array to be sent to the event service.
     * This value cannot be null.
     */
}
```

```

* @return The global instance IDs of the events that were built.
* @exception GenericEmitterException If an error occurs during emitter processing.
* @exception IllegalArgumentException If the events parameter is null.
*/
public String[] sendEvents(GenericEvent events[]) throws
    GenericEventException;
/**
* Causes the emitter to release all resources that are owned by this
* object and its dependents.
* Subsequent calls to this method have no effect.
*
* @throws GenericEmitterException If the emitter does release the
* held resources.
* resources.
* @throws GenericEventException If any other error occurs when releasing resources.
*/
public void close() throws
    GenericEventException;
}

```

Configuring audit event factories for security auditing

The audit event factory collects the data associated with the auditable security events and builds the audit data object. The object is then sent to the audit service provider to be formatted and recorded to a specified repository.

Before you begin

Before configuring an event factory, enable global security in your environment. An event type filter and an audit service provider need to be created before completing these steps

About this task

1. Click **Security** → **Security Auditing** → **Audit event factory configurations** → **New**.
2. Enter the unique name that should be associated with this Audit event factory configuration in the Name field.
3. Select either **IBM audit event factory** or **Third party event factory**.
 - a. Enter the Third party audit event factory class name. This step is only required if a Third party event factory is being created.
4. Select the appropriate audit service provider implementation from the Audit service provider dropdown menu,
5. Select the event type filter configuration to be used by this audit event factory. The Filters list consists of a list of the event type filter configurations that have been created and are currently enabled.
 - a. Select the event type filters that should be used from the Selectable filter list.
 - b. Click **Add >>** to add the selected event type filter configurations to the Enabled filter lists.
6. Enter any Custom properties that need to be included with this audit event factory configuration. Custom properties are only available for Third party event factory implementations.
7. Click **Apply**.

Results

After successful completion of these steps, you will have an event factory that can be used to gather auditable event data.

What to do next

After configuring an audit event factory, you can optionally protect your data by configuring the security auditing subsystem to sign and encrypt your audit logs.

Audit event factory configuration collection

The Audit event factory configuration panel displays a list of all currently configured audit event factory implementations. This panel allows a user with the auditor role to manage their configured audit event factories. This includes the ability to configure a new implementation, which is done using the **New** button on this panel.

To view this administrative console page, click **Security > Security Auditing > Audit event factory configuration**.

Name

The Name field specifies the unique name associated with the audit event factory configuration.

Type

The Type field specifies this audit event factory configuration as either an IBM audit event factory or a Third party audit event factory.

Class name

The Class name field specifies the class that is being implemented in an audit event factory configuration.

The class name is `com.ibm.ws.security.audit.AuditEventFactoryImpl` for an IBM event factory. For a Third party audit event factory, the class name is the class specified in the Third party audit event factory class name field.

Audit event factory settings

The Audit event factory settings panel displays the details of a specific audit event factory. The auditor uses this panel to manage and create audit event factory configurations.

To view this administrative console page, click on of the following paths:

- **Security > Security Auditing > Audit event factory configuration > *audit_event_factory_configuration_name***.
- **Security > Security Auditing > Audit event factory configuration > New**.

Name

Specifies the unique name associated with the audit event factory configuration.

Type

Specifies this audit event factory configuration as either an IBM audit event factory or a Third party audit event factory. This field does not appear on the panel during the creation of a new audit event factory. It is included when viewing or modifying an existing audit event factory.

IBM audit event factory

Specifies that the Type field of this audit event factory is IBM audit event factory. This check box only appears on the panel during the creation of a new audit event factory. This check box is selected by default when creating a new audit event factory.

Third party audit event factory

Specifies that the Type field of this audit event factory is Third party audit event factory. This check box only appears on the panel during the creation of a new audit event factory. This check box is not selected by default when creating a new audit event factory.

- The Third party audit event factory class name field is active when the Third party audit event factory check box is selected. This field represents the class name of the third-party implementation of the Audit Event Factory interface

Class name

Specifies the class that is being implemented in a audit event factory configuration.

Although not specified during creation, the class name is `com.ibm.ws.security.audit.AuditEventFactoryImp` for an IBM event factory.

Audit service provider

Specifies where the audit data objects gathered by this audit event factory will be sent.

Selectable filters

Specifies the filters that are currently available to be used for an implementation. To enable a filter for an implementation, select the filter from the Selectable filter list and then click >.

Enabled filters

Specifies the filters that are currently enabled for an implementation. To disable a filter for an implementation, select the filter from the Enabled filter list and then click <.

Custom properties

Specifies properties that the auditor can define to configure the Audit Event Factory implementation. This might be used by third party implementation of the audit event factory interface. Custom properties are not used for the IBM audit event factory implementation.

Each custom property has the following fields:

- Name
- Value

Example: Generic Event Factory Interface

This interface is used for processing generic audit events. Other interfaces can be defined which extend this interface to process specific audit event groupings, such as security events, transaction events, or some other custom grouping.

Generic Event Factory Interface

```
/**
 * GenericEventFactory is the interface that is used to generate audit events.
 * This interface may be extended to generate application specific audit events.
 *
 * One or more GenericEventFactory implementations each with a unique name can be defined in the
 * security configuration and be used by WebSphere Application Server security auditing service.
 * @author IBM Corporation
 * @version WAS 7.0
 * @since WAS 7.0
 */
public interface GenericEventFactory {
    /**
     * The init method allows a GenericEventFactory implementation to
     * initialize its internal auditing configuration using the properties and context object.
     *
     * The properties and context objects are treated as read-only and must not be modified by the
     * GenericEventFactory implementation.
     *
     * @param A String object represents the name of this GenericEventFactory.
     * @param A Map properties object that contains the custom properties that can be defined in the
     * the admin console or by using wsadmin scripting tool.
     * @param A Map object that contains the context that includes cell name, node name, and server name.
     * @exception ProviderFailureException might occur if the audit factory does not initialize
     */
    public void init(String name, Map properties, Map context) throws ProviderFailureException;
    /**
     * The terminate method gracefully quiesces the event factory implementation.
     */
    public void terminate();
    /**
     * The refresh method allows a GenericEventFactory implementation to
     * update its internal auditing configuration using the properties object.
     *
     * The properties object is treated as read-only and must not be modified by the
     * GenericEventFactory implementation.
     *
     * @param A Map object that contains the custom properties
     * @exception ProviderFailureException might occur if the factory does not refresh
     */
    public void refresh(java.util.Map properties) throws ProviderFailureException;
}
```

```

/**
 *
 * The getName method returns the name of this GenericEventFactory.
 *
 * @param None
 * @return a String object represents the name of the GenericEventFactory.
 */
public String getName();
/**
 *
 * The sendEvent method determines whether the specified audit event is generated by this
 * GenericEventFactory.
 *
 * @param a String object represents an audit event
 * @param a OutcomeType object represents the audit outcome value
 * @exception ProviderFailureException might occur if the audit factory does not initialize
 * @return a boolean success/failure
 * @exception ProviderFailureException might occur if the audit factory does not send the event.
 */
public boolean sendEvent(String auditEventType, OutcomeType auditOutcome) throws
ProviderFailureException;
}

```

Protecting your security audit data

The security auditing subsystem allows for protection of your security audit data by increasing the assurance that the audit data has not been tampered or modified outside of the auditing facility. This option also protects the confidentiality of the data. The audit data is protected by encrypting and signing the recording data.

Before you begin

Note: Signing and encrypting your audit data is only available for data created using the default binary log audit service provider. If you are using the SMF emitter or a 3rd party emitter you will not be able to sign or encrypt your data.

Before configuring protection for your security audit data, enable global security and security auditing in your environment. You must be assigned the auditor role to complete the task of protecting your audit data. You will also need the administrator role to configure your audit data to be signed.

About this task

The practice of auditing requires assurances that your audit data is accurate and uncompromised. Your audit data has the option to be encrypted, signed, or encrypted and signed. You can protect your audit data using these options to provide assurances that you data is only viewed by authorized users and can not untraceably be modified . To protect the validity of your security auditing functionality, complete the following steps:

1. “Encrypting your security audit records” on page 864 Audit logs can be encrypted to ensure your audit data is protected. The audit logs will be encrypted using a certificate that is saved to a keystore in the `audit.xml` file. By encrypting your audit records, only users with the password to the keystore will be able to view or update the audit logs.
2. “Signing your security audit records” on page 865 Audit logs can be signed to ensure the integrity of your audit data. By signing your audit records, you ensure any modifications of the audit logs can be traced.

Results

After completing these steps your data will be signed, encrypted or signed and encrypted to provide assurances that the data is accurate and confidential.

What to do next

After protecting your data, you can configure notifications to ensure you are notified if a problem with the security auditing subsystems occurs that prevents security events from being recorded.

Encrypting your security audit records

Audit logs can be encrypted to ensure your audit data is protected. By encrypting your audit records, only users with access to the encrypting certificate will be able to view the audit logs.

Before you begin

Note: Encrypting audit data is only available for data created using the default audit service provider. If you are using the SMF emitter or a 3rd party emitter you will not be able to encrypt your data.

Before configuring your security audit records to be encrypted, enable global security and security auditing in your environment. You must be assigned the auditor role to encrypt your security auditing records. If you are using a certificate stored in the security.xml file, you also require the administrator role to complete this task.

About this task

1. Click **Security** → **Security Auditing** → **Audit record encryption configuration**.
2. Select the Enable encryption check box to specify that your audit records should be encrypted. All other fields on this panel will be unavailable until this check box has been selected.
3. Select the keystore that contains the encrypting certificate from the dropdown menu or click **New** to create a new certificate in an existing keystore. Use the following steps if you are creating a new certificate:
 - a. Enter the name of the keystore in the Name field.
 - b. Enter the path to the keystore file in the Path field.
 - c. Enter the password to be associated with the keystore in the Password field.
 - d. Confirm the password associated with the keystore by retyping the password in the Confirm password field.
 - e. Select the keystore type from the Type dropdown list. The default value of the Type dropdown list is PKCS12.
4. If you are using an existing certificate to encrypt your audit records, ensure **Certificate in keystore** is selected and specify the intended certificate in the Certificate alias dropdown menu.
5. If you are generating a new certificate to encrypt your audit records, select **Create a new certificate in the selected keystore** and follow these steps:
 - a. Enter the name of your new certificate in the Certificate alias field.
 - b. Select either Automatically generate certificate or Import a certificate. The certificate used to encrypt the data in the audit log files can either be created or imported. If you selected to generate a certificate, then skip to the last step on this page. If you selected to import a certificate, then continue on with step c.
 - c. Enter the name of the keystore file in the Key file name field.
 - d. Enter the path to the keystore file in the Path field.
 - e. Select the keystore type from the Type dropdown list. The default value of the Type dropdown list is PKCS12.
 - f. Enter the password associated with the keystore in the Key File password field.
 - g. Click **Get key file aliases** to populate the Certificate alias to import dropdown menu.
 - h. Select the certificate to be imported from the Certificate alias to import dropdown menu.
6. Click **OK**.

Results

After completing these steps, your audit logs will be encrypted to ensure only authorized users can view the content of your audit log files.

What to do next

After you have finished configuring your audit logs to be encrypted, you can ensure the data integrity of your audit logs by configuring the audit subsystem to sign your audit records.

Signing your security audit records

Audit logs can be signed to ensure the integrity of your audit data. By signing your audit records, modifications of the audit logs can be traced.

Before you begin

Note: Signing audit data is only available for data created using the default audit service provider. If you are using the SMF emitter or a 3rd party emitter you will not be able to sign your data.

Before configuring your security audit records to be signed, enable global security and security auditing in your environment. You must be assigned the auditor role and the administrator role to configure audit record signing.

About this task

1. Click **Security** → **Security Auditing** → **Audit record signing configuration**.
2. Select the Enable signing check box to specify that your audit records should be signed. All other fields on this panel will be unavailable until this check box has been selected.
3. Select the keystore that contains the signing certificate from the Managed keystore containing the signing certificate dropdown menu.
4. If you are using an existing certificate to sign your audit records, ensure Certificate in keystore is selected and specify the intended certificate in the Certificate alias dropdown menu.
5. If you are generating a new certificate to sign your audit records, select Create a new certificate in the selected keystore and follow these steps:
 - a. Enter the name of your new certificate in the Certificate alias field.
 - b. Select one of the following options: Import the encryption certificate, Automatically generate certificate or Import a certificate. The certificate used to encrypt the data in the audit log files can either be created or imported.
 - If you selected Import the encryption certificate, then you will use the encryption certificate to also sign your audit records. Skip to the last step on this page to complete this configuration.
 - If you selected to generate a certificate, then skip to the last step on this page to complete this configuration.
 - If you selected to import a certificate from an existing keystore, then continue on with step c.
 - c. Enter the name of the keystore file in the Key file name field.
 - d. Enter the path to the keystore file in the Path field.
 - e. Select the keystore type from the Type dropdown list. The default value of the Type dropdown list is PKCS12.
 - f. Enter the password associated with the keystore in the Key File password field.
 - g. Click **Get key file aliases** to populate the Certificate alias to import dropdown menu.
 - h. Select the certificate to be imported from the Certificate alias to import dropdown menu.
6. Click **OK**.

Results

After you have completed these steps, your audit logs will be digitally signed to ensure the integrity of the data.

What to do next

After you have finished configuring your audit logs to be signed, you can ensure the confidentiality of your audit logs by configuring the audit subsystem to encrypt your audit records.

Audit encryption keystores and certificates collection

The Audit encryption keystores and certificates panel allows the auditor to manage the keystores and certificates used for audit encryption.

To view this administrative console page, click **Security** → **Security Auditing** → **Audit encryption keystores and certificates**.

Name

Specifies the unique name of the keystores used for storing the encryption certificate.

Path

Specifies the path to the listed keystore file.

The path to the keystore file can be listed using environment variables, `${PROFILE_ROOT}`, or with a fully qualified path.

Audit record encryption configuration settings

Use this page to enable encryption for your audit records. Encrypting your audit records ensures only a user given access to the certificate used for encryption is allowed to view the audit records.

To view this administrative console page, click **Security** > **Security auditing** > **Audit record encryption configuration**. If **Enable encryption** is not selected, then all of the other fields on this panel will be disabled. Encryption is not enabled by default.

Enable encryption

Specifies whether your audit records will be encrypted. This check box is not selected by default.

Audit keystore containing the encryption certificate

Specifies the audit keystore specified to store the encryption certificate.

A new keystore can be created by clicking on the **New...** button.

Certificate in keystore

Specifies an existing certificate will be used from the keystore specified in the Audit keystore containing the encryption certificate field. This field is selected by default. If a keystore in the security.xml file is used, the administrator role is required.

- Certificate alias

When the **Certificate in keystore** field is selected, the certificate alias dropdown menu displays a list of certificate aliases contained in the keystore defined by the **Audit keystore containing the encryption certificate** field. Select the certificate from the dropdown menu to be used to encrypt your audit records.

Create a new certificate in the selected keystore

Specifies that a new certificate will be created in the keystore defined by the **Audit keystore containing the encryption certificate** field.

- Certificate alias

When the **Create a new certificate in the selected keystore** is selected, the **Certificate alias** field is used to define the name of the certificate to be created in the keystore defined by the **Audit keystore containing the encryption certificate** field.

- Automatically generate certificate

When selected, the **Automatically generate certificate** field specifies that the application server will automatically generate the certificate. This field is selected by default when the **Create a new certificate in the selected keystore** field is selected.

- Import a certificate

When selected, the **Import a certificate** field specifies that an existing self-signed certificate will be imported by the auditor into the keystore and used to encrypt your audit records. This field is not selected by default when the **Create a new certificate in the selected keystore** field is selected. The following fields need to be defined to import an existing certificate.

- The **Key file name** field specifies the keystore filename that contains the certificate to be imported.
- The **Path** field specifies the path to the keystore file that contains the certificate to be imported.
- The **Type** field specifies the type of the keystore file that contains the certificate to be imported.
- The **Key file password** field specifies the password used to access the keystore file that contains the certificate to be imported.
- **Certificate alias to import** field specifies the alias of the certificate to be imported.

Audit record signing configuration settings

Use this page to enable signing for your audit records. Signing audit records ensures tamper-proof recording of the auditable events. Both the auditor and administrator roles are required to configure the signing of your audit data.

To view this administrative console page, click **Security** → **Security auditing** → **Audit record signing configuration**. If **Enable signing** is not selected, then all of the other fields on this panel will be disabled.

Related reference

“Audit record keystore settings” on page 868

The Audit record keystore panel is used by an auditor to define the keystores used for storing the encryption certificate used to encrypt the audit records. Keystores used for auditing are managed outside of other keystores being used on the system to facilitate separation of the authority of the auditor for the authority of the administrator.

“Security Auditing detail” on page 842

The Security auditing subsystem can be enabled and configured from this panel, by users assigned the auditor role.

“Audit record encryption configuration settings” on page 866

Use this page to enable encryption for your audit records. Encrypting your audit records ensures only a user given access to the certificate used for encryption is allowed to view the audit records.

“Audit encryption keystores and certificates collection” on page 866

The Audit encryption keystores and certificates panel allows the auditor to manage the keystores and certificates used for audit encryption.

Enable signing

Specifies whether your audit records will be encrypted. This check box is not selected by default.

Managed keystore containing the signing certificate

Specifies the keystore used to store the signing certificate.

Certificate in keystore

Specifies an existing certificate will be used from the keystore specified in the **Managed keystore containing the signing certificate** field. This field is selected by default.

- Certificate alias

When the **Certificate in keystore** field is selected, the **Certificate alias** dropdown menu displays a list of certificate aliases contained in the keystore defined by the **Managed keystore containing the signing certificate** field. Select the certificate from the dropdown menu to be used to sign your audit records.

Create a new certificate in the selected keystore

Specifies that a new certificate will be created in the keystore defined by the **Managed keystore containing the signing certificate** field.

- Certificate alias

When the **Create a new certificate in the selected keystore** is selected, the **Certificate alias** field is used to define the name of the certificate to be created in the keystore defined by the **Audit keystore containing the encryption certificate** field.

- Import the encryption certificate

Specifies the certificate used for encryption will be imported into the signing keystore file and used for signing.

- Automatically generate certificate

Specifies the application server will automatically generate the certificate. This field is selected by default when the **Create a new certificate in the selected keystore** field is selected.

- Import a certificate

Specifies an existing self-signed certificate will be imported by the auditor into the keystore and used to encrypt your audit records. This field is not selected by default when the **Create a new certificate in the selected keystore** field is selected. The following fields need to be defined to import an existing certificate.

- The **Key file name** field specifies the keystore filename that contains the certificate to be imported.
- The **Path** field specifies the path to the keystore file that contains the certificate to be imported.
- The **Type** field specifies the type of the keystore file that contains the certificate to be imported.
- The **Key file password** field specifies the password used to access the keystore file that contains the certificate to be imported.
- **Certificate alias to import** field specifies the alias of the certificate to be imported.

Audit record keystore settings

The Audit record keystore panel is used by an auditor to define the keystores used for storing the encryption certificate used to encrypt the audit records. Keystores used for auditing are managed outside of other keystores being used on the system to facilitate separation of the authority of the auditor for the authority of the administrator.

To view this administrative console page, click one of the following paths:

- **Security > Security Auditing > Audit encryption keystores and certificates > *keystore_name*.**
- **Security > Security Auditing > Audit encryption keystores and certificates > New.**
- **Security > Security Auditing > Audit record encryption configuration > New**

Related reference

“Security Auditing detail” on page 842

The Security auditing subsystem can be enabled and configured from this panel, by users assigned the auditor role.

“Audit record signing configuration settings” on page 867

Use this page to enable signing for your audit records. Signing audit records ensures tamper-proof recording of the auditable events. Both the auditor and administrator roles are required to configure the signing of your audit data.

“Audit record encryption configuration settings” on page 866

Use this page to enable encryption for your audit records. Encrypting your audit records ensures only a user given access to the certificate used for encryption is allowed to view the audit records.

“Audit encryption keystores and certificates collection” on page 866

The Audit encryption keystores and certificates panel allows the auditor to manage the keystores and certificates used for audit encryption.

Name

The Name field specifies the unique name for the keystore. This is a required field.

Path

Specifies the path where the keystore file is located. This is a required field.

Password

Specifies the password to be used for this keystore. This is a required field.

Confirm Password

Specifies confirmation of the value provided in the Password field. This is a required field.

Type

The Type field specifies the type of the keystore. The Type dropdown menu has the following options for defining the keystore type:

- JCEKS
- CMSKS
- PKCS12 - The default value for the Type field is PKCS12.
- Cryptographic Token Device (PKCS11)
- JKS
- PKCS12JarSigner

Using the audit reader

The audit reader is a utility that can be used to read the binary audit logs generated by the default binary emitter implementation. The audit reader parses the audit log to generate an HTML report. The audit reader is invoked using wsadmin commands and is not accessible using the administrative console.

Before you begin

The audit reader can only be used to parse log files that are created by the default audit service provider. Logs created by a third-party emitter can not be parsed by the audit reader.

About this task

Your audit logs might be encrypted, signed, encrypted and signed or neither encrypted nor signed. The audit reader is able to parse any of these combinations to generate an HTML report. If the audit log file is

encrypted, the password of the keystore storing the certificate used to encrypt the log must be provided. The `showAuditLogEncryptionInfo wsadmin` command can be used to get information to determine which keystore was used to sign the audit log.

Depending on the selections you made in your audit service provider configuration, the size of the audit logs can become large enough to make them cumbersome to review. What data has been recorded into your log is dependant on the event type filters you are using and whether you specified to use verbose logging. Options are provided for you to further limit the data included in the HTML report that is generated by the audit reader to a subset that you specify. The audit reader can be used to parse the same data multiple times to generate separate reports for your different requirements.

By default, all event types, outcome types, timestamps, and sequence numbers will be gathered from the Binary Audit log and generated into a report. The ability to specify only specific event types, only specific sequence numbers, only records with specific timestamps, as well as specific outcome types is provided. A sequence number is a unique identifier assigned to each audit record. Options exist to limit which events, outcomes, and sequence numbers are included in the report.

The report type controls what data is reported for each audit record in the log file. The default report type includes the follow data for each audit record:

- `creationTime`
- `action`
- `progName`
- `registryType`
- `domain`
- `realm`
- `remoteAddr`
- `remotePort`
- `remoteHost`
- `resourceName`
- `resourceType`
- `resourceUniqueld`

The complete report type generates a report based on all the data that was logged for the selected audit records. The complete report type includes all the data that is included by the default report type and all the additional datapoints that were logged for these audit records. The additional available datapoints for an audit record varies depending on the event type it represents.

A custom report type is also included. Use the custom report type to specify only the datapoints that you want generated in the report. A report may be generated based on the following criteria:

- all or specific event types
- all or specific outcome types
- all or a specific sequence number range
- all or a specific timestamp range

Run the `binaryAuditLogReader wsadmin` command to use the audit reader to generate a log report. See “`AuditReaderCommands` command group for the `AdminTask` object” on page 1152

Results

After you complete these steps, you will generated an HTML report containing the data specific to your requirement.

Example: Audit Event Outcome Codes

In a binary audit log or the output of the audit reader tool, audit event outcomes are expressed with a numeric code. Use this table to associate the audit event outcome code in the binary audit logs to a generic error messages.

Event Outcome Codes

Outcome Reason Code	Description
0	An error occurred while parsing the certificate.
1	The security context does not exist for the thread.
2	There is conflicting session evidence.
3	The session has been rejected.
4	The token has expired.
5	Successful authentication has occurred.
6	Successful authentication for accessing a resource has occurred.
7	Successful authentication occurred while mapping a user.
8	Successful authorization has occurred.
9	Login termination was successful.
10	Invalid evidence exists.
11	There was a GSS formatting error.
12	Credentials were unauthenticated.
13	Authentication failed.
14	An invalid resource was accessed.
15	Authentication was denied.
16	Authorization was denied.
17	Access was denied because of an authentication failure.
18	Authorization was excluded.
19	Authorization was excluded because of access without proper security role.
20	An unsupported authentication mechanism was used.
21	An authentication redirect occurred.
22	The context does not exist.
23	A TAI challenge occurred.
24	A TAI validation was not successful.
25	A TAI mapping was not successful.
26	A provider failure occurred.
27	A SSO token validation was not successful.
28	An invalid user id or password was provided.
29	A send login form
30	An invalid configuration exists.
31	An user id or password is missing.
32	Failure occurred for an unknown reason.
33	The account was disabled because of retry violations.
34	The account was locked out because of retry violations.
35	The account was locked out because the maximum number of unsuccessful login attempts has occurred.
36	The account is disabled.
37	The account has expired.
38	The account is unlocked.
39	The maximum inactive time permitted for the account has elapsed.
40	The password has expired.
41	The minimum interval for a password change has unexpired.
42	The maximum interval permitted before a password must be changes has elapsed.
43	An authentication failure has occurred.

44	An invalid user name was provided.
45	A pin is required.
46	This outcome code is not used in this release.
47	A user mapping did not occur successfully.
48	A certificate failure occurred.
49	A policy violation has occurred.
50	A policy violation has occurred because of the time of day.
51	The policy allows access.
52	A policy violation has occurred because the maximum number of unsuccessful login attempts has been reached.
53	A user name mismatch has occurred.
54	An invalid user password was provided.
55	A token signature violation has occurred.
56	The token is not yet valid.
57	The token is not supported.
58	The token is not in a valid format.
59	A credential mapping failure occurred.
60	The delegate is not authorized.
61	Access to a resource is unauthorized because of an authorization.
62	Access to a resource is unauthorized because of a time of day policy.
63	Access to a resource is unauthorized.
64	Access to a resource is unauthorized because of quality of protection.
65	Access to a resource is unauthorized because of an authorization level.
66	Access to a resource is unauthorized because reauthentication is required.
67	A password error has occurred because it does not meet password standards: minimum alphabetic characters required.
68	A password error has occurred because it does not meet password standards: minimum alphanumeric characters required.
69	A password error has occurred because it does not meet password standards: minimum numeric characters required.
70	A password error has occurred because it does not meet password standards: minimum alphabetic low case characters required.
71	A password error has occurred because it does not meet password standards: minimum alphabetic upper case characters required.
72	A password error has occurred because it does not meet password standards: minimum special characters required.
73	A password error has occurred because it does not meet password standards: maximum repeated characters exceeded.
74	A password error has occurred because it does not meet password standards: contains user name
75	A password error has occurred because it does not meet password standards: reused password.
76	A password error has occurred because it does not meet password standards: contains previous password.
77	A password error has occurred because it does not meet password standards: violations in number of characters.
78	A password error has occurred because it does not meet password standards: first or last characters are numeric.
79	An illegal form login configuration exists.
80	Access is denied because of a incorrect URI.
81	Start was successful
82	Stop was successful.
83	The audit subsystem has been stopped.
84	The audit subsystem has successfully been enabled.
85	The audit subsystem has had a successful policy change.
86	Delegation was successful.

87	Delegation was not successful.
88	The audit subsystem has successfully been disabled.
89	An audit subsystem has occurred because a security header is missing.
90	An audit timestamp has been confirmed.
91	A bad audit timestamp has occurred.
92	Audit confidentially has been confirmed
93	Audit confidentially cannot be confirmed.
94	An audit decryption error has occurred.

Chapter 10. Configuring security with scripting

You can configure security with scripting and the wsadmin tool.

Before you begin

Before starting this task, the wsadmin tool must be running. Read about Starting the wsadmin scripting client article for more information.

About this task

If you enable security for an application server cell, supply authentication information to communicate with servers. The `sas.client.props` and the `soap.client.props` files are located in the following properties directory for each application server profile:

- `profile_root/properties`
- The nature of the properties file updates required for running in secure mode depend on whether you connect with a Remote Method Invocation (RMI) connector, a JSR160RMI connector, an Inter-Process Communications (IPC) or a SOAP connector:
 - If you use a Remote Method Invocation (RMI) connector or a JSR160RMI connector, set the following properties in the `sas.client.props` file with the appropriate values:

```
com.ibm.CORBA.loginUserId=  
com.ibm.CORBA.loginPassword=
```

Also, set the following property:

```
com.ibm.CORBA.loginSource=properties
```

The default value for this property is `prompt` in the `sas.client.props` file. If you leave the default value, then a dialog box is displayed with a password prompt. If the script is running unattended, then the system stops.

- If you use a SOAP connector, set the following properties in the `soap.client.props` file with the appropriate values:

```
com.ibm.SOAP.securityEnabled=true  
com.ibm.SOAP.loginUserId=  
com.ibm.SOAP.loginPassword=
```

Optionally, set the following property:

```
com.ibm.SOAP.loginSource=none
```

The default value for this property is `prompt` in the `soap.client.props` file. If you leave the default value, a dialog box is displayed with a password prompt. If the script is running unattended, then the system stops.

- If you use an IPC connector, set the following properties in the `ipc.client.props` file with the appropriate values:

```
com.ibm.IPC.loginUserId=  
com.ibm.IPC.loginPassword=
```

Optionally, set the following property:

```
com.ibm.IPC.loginSource=prompt
```

The default value for this property is `prompt` in the `soap.client.props` file. If you leave the default value, a dialog box appears with a password prompt. If the script is running unattended, it appears to hang.

- Specify user and password information. Choose one of the following methods:
 - Specify user name and password on a command line, using the **-user** and **-password** commands, as the following examples demonstrate:

```
wsadmin -conntype JSR160RMI -port 2809 -user u1 -password secret1
```

- Specify user name and password in the `sas.client.props` file for an RMI connector, the `ipc.client.props` file for the IPC connector, or the `soap.client.props` file for a SOAP connector.

If you specify user and password information on a command line and in the `sas.client.props` file or the `soap.client.props` file, the command line information overrides the information in the props file.

AIX

HP-UX

Linux

Solaris

The use of `-password` option may result in security exposure as the password information becomes visible to the system status program such as `ps` command which can be invoked by other user to display all the running processes. Do not use this option if security exposure is a concern. Instead, specify user and password information in the `soap.client.props` file for the SOAP connector, the `sas.client.props` file for the JSR160RMI connector or the Remote Method Invocation (RMI) connector, or the `ipc.client.props` file for the IPC connector. The `soap.client.props`, `sas.client.props`, and `ipc.client.props` files are located in the properties directory of your profile.

Enabling and disabling security using scripting

You can use scripting to enable or disable application security, global security, administrative security based on the LocalOS registry, and authentication mechanisms.

Before you begin

Before starting this task, the `wsadmin` tool must be running. See the [Starting the wsadmin scripting client](#) article for more information.

About this task

The default profile sets up procedures so that you can enable and disable administrative security based on LocalOS registry.

- Use the `isAppEnabled` command to determine if application security is enabled or disabled, as the following example demonstrates:

- Using Jacl:

```
$AdminTask isAppSecurityEnabled {}
```

- Using Jython:

```
AdminTask.isAppSecurityEnabled()
```

This command returns a value of `true` if `appEnabled` is set to `true`. Otherwise, returns a value of `false`.

- Use the `isGlobalSecurityEnabled` command to determine if administrative security is enabled or disabled, as the following example demonstrates:

- Using Jacl:

```
$AdminTask isGlobalSecurityEnabled {}
```

- Using Jython:

```
AdminTask.isGlobalSecurityEnabled()
```

Returns a value of `true` if `enabled` is set to `true`. Otherwise, returns a value of `false`.

- Use the `setGlobalSecurity` command to set administrative security based on the passed in value, as the following example demonstrates:

- Using Jacl:

```
$AdminTask setGlobalSecurity {-enabled true}
```

- Using Jython:

```
AdminTask.setGlobalSecurity ('[-enabled true]')
```

Returns a value of `true` if the `enabled` field in the WCCM security model is successfully updated. Otherwise, returns a value of `false`.

- Use the **help** command to find out the arguments that you need to provide with this call, as the following example demonstrates:
 - Using Jacl:


```
securityon help
```

 Example output:


```
Syntax: securityon user password
```
 - Using Jython:


```
securityon()
```

 Example output:


```
Syntax: securityon(user, password)
```
- Enable administrative security based on the LocalOS registry, as the following example demonstrates:
 - Using Jacl:


```
securityon user1 password1
```
 - Using Jython:


```
securityon('user1', 'password1')
```
- Disable administrative security based on the LocalOS registry, as the following example demonstrates:
 - Using Jacl:


```
securityoff
```
 - Using Jython:


```
securityoff()
```
- Enable and disable LTPA and Kerberos authentication.

Use the `setActiveAuthMechanism` command to set Kerberos as the authentication mechanism in the security configuration, as the following example demonstrates:

```
AdminTask.setActiveAuthMechanism('-authMechanismType KRB5')
```

Use the `setActiveAuthMechanism` command to set LTPA as the authentication mechanism in the security configuration, as the following example demonstrates:

```
AdminTask.setActiveAuthMechanism('-authMechanismType LTPA')
```

Additionally, there are sample scripts located in the `<WAS_ROOT>/bin` directory on how to enable and disable LTPA authentication. The scripts are:

 - `LTPA_LDAPSecurityProcs.py` (python script)
 - `LTPA_LDAPSecurityProcs.jacl` (jacl script)

Note: The scripts hard code the type of LDAP server and base distinguished name (baseDN). The LDAP server type is hardcoded as `IBM_DIRECTORY_SERVER` and the baseDN is hardcoded as `o=ibm,cn=us`.

Enabling and disabling Java 2 security using scripting

You can enable or disable Java 2 security with scripting and the `wsadmin` tool.

About this task

There are two ways to enable or disable Java 2 security. You can use the commands for the `AdminConfig` object, or you can use the `setAdminActiveSecuritySettings` command for the `AdminTask` object.

1. Use the `setAdminActiveSecuritySettings` command for the `AdminTask` object to enable or disable Java 2 security.
 - a. Launch the `wsadmin` scripting tool using the Jython scripting language.
 - b. Use the `getActiveSecuritySettings` command to display the current security settings, including custom properties for global security, as the following example demonstrates:

- - Using Jacl:


```
$AdminTask getActiveSecuritySettings
```
 - Using Jython:


```
AdminTask.getActiveSecuritySettings()
```
- c. Use the `setActiveSecuritySettings` command to enable or disable Java 2 security. The following examples enable Java 2 security:

- - Using Jacl:


```
$AdminTask setAdminActiveSecuritySettings {-enforceJava2Security true}
```
 - Using Jython:


```
AdminTask.setAdminActiveSecuritySettings('-enforceJava2Security true')
```

The following examples disable Java 2 security:

- - Using Jacl:


```
$AdminTask setAdminActiveSecuritySettings {-enforceJava2Security false}
```
 - Using Jython:


```
AdminTask.setAdminActiveSecuritySettings('-enforceJava2Security false')
```
- d. Save the configuration changes.
- e. Synchronize the node.

Use the `syncActiveNode` or `syncNode` scripts in the `AdminNodeManagement` script library to propagate the configuration changes to node or nodes.

 - Use the `syncActiveNodes` script to propagate the changes to each node in the cell, as the following example demonstrates:


```
AdminNodeManagement.syncActiveNodes()
```
 - Use the `syncNode` script to propagate the changes to a specific node, as the following example demonstrates:


```
AdminNodeManagement.syncNode("myNode")
```

2. Use the `AdminConfig` object to enable Java 2 security.

- a. Launch the `wsadmin` scripting tool using the Jython scripting language.
- b. Identify the security configuration object and assign it to the security variable, as the following example demonstrates:
 - Using Jacl:


```
set security [$AdminConfig list Security]
```
 - Using Jython:


```
security = AdminConfig.list('Security')
print security
```

Example output:

```
(cells/mycell|security.xml#Security_1)
```
- c. Modify the `enforceJava2Security` attribute to enable or disable Java 2 security, as the following examples demonstrate:
 - To enable Java 2 security:
 - Using Jacl:


```
$AdminConfig modify $security {{enforceJava2Security true}}
```
 - Using Jython:


```
AdminConfig.modify(security, [['enforceJava2Security', 'true']])
```
 - To disable Java 2 security:

- Using Jacl:


```
$AdminConfig modify $security {{enforceJava2Security false}}
```
- Using Jython:


```
AdminConfig.modify(security, [['enforceJava2Security', 'false']])
```

- d. Save the configuration changes.
- e. Synchronize the node.

Use the `syncActiveNode` or `syncNode` scripts in the `AdminNodeManagement` script library to propagate the configuration changes to node or nodes.

- Use the `syncActiveNodes` script to propagate the changes to each node in the cell, as the following example demonstrates:

```
AdminNodeManagement.syncActiveNodes()
```

- Use the `syncNode` script to propagate the changes to a specific node, as the following example demonstrates:

```
AdminNodeManagement.syncNode("myNode")
```

Configuring multiple security domains using scripting

You can customize your security configuration at the cell, sever, or cluster level by configuring multiple security domains.

Before you begin

Users assigned to the administrator role can configure security domains. Verify that you have the appropriate administrative role before configuring security domains. Also, enable global security in your environment before configuring multiple security domains.

About this task

You can create multiple security domains to customize your security configuration. Use multiple security domains to achieve the following goals:

- Configure different security attributes for administrative and user applications within a cell
 - Consolidate server configurations by managing different security configurations within a cell
 - Restrict access between applications with different user registries, or configure trust relationships between applications to support communication across registries
1. Create a security domain. Create multiple security domains in your configuration. By creating multiple security domains, you can configure different security attributes for administrative and user applications within a cell environment.
 2. Assign the security domain to one or a set of resources or scopes. Assign management resources to security domains. Set management resources to your security domains to customize your security configuration for a cell, server, or cluster.
 3. Customize your security configuration by specifying attributes for your security domain. See the following examples of security attributes:
 - User registries to validate user credentials
 - Authorization for validating access to resources
 - Trust association interceptor (TAI) to authenticate a Web user using a reverse proxy server
 - Application and system JAAS login configurations
 - LTPA timeout settings
 - Application security enablement to provide application isolation and requirements for authenticating application users

- Java 2 Security to increase overall system integrity by checking for permissions before allowing access to certain protected system resources
- Remote Method Invocation over Internet Inter-ORB Protocol (RMI/IIOP) to invoke Web services through remote procedure calls
- Custom properties

Configuring security domains using scripting

Use this topic to create multiple security domains in your configuration. By creating multiple security domains, you can configure different security attributes for administrative and user applications within a cell environment.

Before you begin

You must have the administrator role to configure security domains. Also, enable global security in your environment before configuring multiple security domains.

About this task

You can create multiple security domains to customize your security configuration. Use multiple security domains to achieve the following goals:

- Configure different security attributes for administrative and user applications within a cell
- Consolidate server configurations by managing different security configurations within a cell
- Restrict access between applications with different user registries, or configure trust relationships between applications to support communication across registries

Use the following steps to create a new security domain with the wsadmin tool:

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Create a security domain.

To create a security domain, you can create a new security domain, copy an existing security domain, or copy the existing global security configuration.

- Use the `createSecurityDomain` command to create the `domain-security.xml` and `domain-security-map.xml` security files in the `profile_root/config/waspolicies/default/security_configuration_name` directory. No configuration data is added to the `domain-security.xml` file. Use the following Jython command example to create a security domain:

```
AdminTask.createSecurityDomain('-securityDomainName securityDomain1 -securityDomainDescription "handles user applications"')
```

The command returns the object name of the security domain that was created, as the following example output demonstrates:

```
'waspolicies/default/securitydomains/mydomain:domain-security.xml#AppSecurity_1183132319126'
```

- Use the `copySecurityDomain` command to create a new security domain with the attributes of an existing security domain. If the security configuration of the existing domain has an active user registry defined, then a new realm name for that registry must be used in the new security configuration. If a realm name is not specified with the `copySecurityDomain` command, then the command assigns a name. Specify the following parameters to copy an existing security domain:

Parameter	Description
-securityDomainName	Specifies the name of the new security domain to create (String, required)
-copyFromSecurityDomainName	Specifies the name of the existing security domain to copy (String, required)
-realmName	Specifies the name of the realm in the security domain to create. The system assigns the realm name to the active user registry in the security domain (String, optional)
-securityDomainDescription	Specifies a description for the security domain to create (String, optional)

Use the following Jython command to copy an existing security domain:

```
AdminTask.copySecurityDomain('-securityDomainName copyOfDomain1 -copyFromSecurityDomainName securityDomain1')
```

The command returns the object name of the new security domain, as the following example output demonstrates:

```
'waspolicies/default/securitydomains/copyOfDomain1:domain-security.xml#AppSecurity_1183132319186'
```

- Use the `copySecurityDomainFromGlobalSecurity` command to create a security domain by copying the global security configuration. If the global security configuration has an active user registry defined, then a new realm name for that registry must be used for the security domain to create. If you do not specify a realm name, then the command assigns a name. Specify the following parameters to copy the global security configuration:

Parameter	Description
-securityDomainName	Specifies the name of the new security domain to create (String, required)
-realmName	Specifies the name of the realm in the security domain to create. The system assigns the realm name to the active user registry in the security domain (String, optional)
-securityDomainDescription	Specifies a description for the security domain to create (String, optional)

Use the following Jython command to copy the global security configuration:

```
AdminTask.copySecurityDomainFromGlobalSecurity('-securityDomainName G$copy')
```

The command returns the object name of the new security domain, as the following example output demonstrates:

```
'waspolicies/default/securitydomains/copyOfDomain1:domain-security.xml#AppSecurity_1183132319186'
```

3. Save your configuration changes.

What to do next

Use the `wsadmin` tool to map a scope to your security domain. Additionally, you can configure security artifacts in the newly created domain, by:

- configuring user registries.
- enabling application and Java EE security.
- setting Lightweight Third-Party Authentication (LTPA) timeout.
- configuring System and Application Java™ Authentication and Authorization Service (JAAS) login.
- configuring Java 2 Connector (J2C) authorization data.
- configuring Remote Method Invocation over Internet Inter-ORB Protocol (RMI/IIOP) security.

Configuring local operating system user registries using scripting

Use this topic to configure user registries for global security and security domain configurations using the `wsadmin` tool. You can define user registries at the global level and for multiple security domains.

Before you begin

You must meet the following requirements before configuring local operating system user registries:

- You must have the administrator or new admin role.
- Enable global security in your environment.
- To configure local operating system user registries for multiple security domains, you must configure at least one security domain.

About this task

Configure local operating system user registries to support use of the authentication mechanism with the user accounts database of the local operating system. You can specify local operating system user registries at the global level and at the security domain.

When you configure a user registry in the global security configuration, the administrator does not specify a realm name for the user registry. The system determines the realm name from the security runtime. The system typically specifies the hostname for local operating system registries.

In security domains, you can configure a different realm for a user registry configuration. For example, you can configure two registries that use the same LDAP server listening on the same port, but use different base distinguished names (baseDN). This allows the configuration to serve different sets of users and groups. To use this type of scenario, you must specify a realm name for each user registry configured for a domain. Because there can be multiple realms in your configuration, you can also specify a list of trusted realms. This allows communication between applications that use different realms.

Use the following steps to configure local operating system user registries for your global security configuration and for multiple security domains:

- Configure local operating system registries for global security configurations.
 1. Use the `configureAdminLocalOSUserRegistry` command and the following optional parameters to configure a local operating system user registry in your global security configuration:

Parameter	Description	Data type
<code>-autoGenerateServerId</code>	Specifies whether to automatically generate the server identity to use for internal process communication. To set a specific server identity, specify the <code>-serverId</code> parameter.	Boolean
<code>-serverId</code>	Specifies the user identity in the repository to use for internal process communication.	String
<code>-serverIdPassword</code>	Specifies the password that corresponds to the user identity.	String
<code>-primaryAdminId</code>	Specifies the name of the user with administrative privileges as defined in the registry. This parameter does not apply to security configurations. The user name must exist in the user registry repository.	String
<code>-customProperties</code>	Specifies a list of attribute and value pairs to store as custom properties on the user registry. Separate each attribute and value pair with a comma character (,), as the following syntax displays: "attribute1=value1","attribute2=value2"	String
<code>-verifyRegistry</code>	Specifies whether to verify the user registry. The default value is <code>true</code> and verification is automatically performed.	Boolean
<code>-ignoreCase</code>	Specifies whether to perform the case-sensitive authorization check. This only applies to the z/OS local operating system user registry.	Boolean

Use the following Jython example command to configure the local operating system registry for global security:

```
AdminTask.configureAdminLocalOSUserRegistry('-autoGenerateServerId true -primaryAdminId gsAdmin')
```

2. Configure the user registry to be the active user registry for the server.

For example, the following Jython command sets the active user registry as the `LocalOSUserRegistry` registry for your global security configuration:

```
AdminTask.setAdminActiveSecuritySettings('-activeUserRegistry LocalOSUserRegistry')
```

3. Save your configuration changes.

- Configure local operating system registries for security domains.

1. Determine the name of the security domain to configure.

Use the `listSecurityDomains` command to list all security domains on the server, as the following Jython example demonstrates:

```
AdminTask.listSecurityDomains()
```

If you want to configure the local operating system registry for a specific server, cluster, or cell, use the `getSecurityDomainForResource` command to display the security domain name for the management scope of interest. The following Jython example displays the name of the security domain configured at the cell-level:

```
AdminTask.getSecurityDomainForResource('-resourceName Cell=:Node=myNode:Server=myServer')
```

For this example, the command returns the following output:

```
domain2
```

2. Configure a local operating system user registry for a security domain.

Use the `configureAppLocalOSUserRegistry` command and the following optional parameters to configure a local operating system user registry:

Parameter	Description	Data type
-securityDomainName	Specifies the unique name that identifies the security domain of interest.	String
-realmName	Specifies the name of the realm of the user registry.	String
-customProperties	Specifies a list of attribute and value pairs to store as custom properties on the user registry object. Separate each attribute and value pair with a comma character (,).	String
-verifyRegistry	Specifies whether to verify the user registry. The default value is true, and verification is automatically performed.	Boolean
-ignoreCase	Specifies whether to perform the case-sensitive authorization check. This only applies to the z/OS local operating system user registry.	Boolean

Use the following Jython command to configure the local operating system user registry for the `domain2` security domain:

```
AdminTask.configureAppLocalOSUserRegistry('-securityDomainName domain2 -realmName domain2Realm')
```

3. Configure the user registry to be the active user registry for the server.

For example, the following Jython command sets the active user registry as the `LocalOSUserRegistry` registry for your security domain configuration:

```
AdminTask.setAppActiveSecuritySettings('-securityDomainName domain2 -activeUserRegistry LocalOSUserRegistry')
```

4. Save your configuration changes.

What to do next

Configuring custom user registries using scripting

Use this topic to configure custom user registries for global security and security domain configurations using the `wsadmin` tool. You can define custom user registries at the global level and for multiple security domains.

Before you begin

You must meet the following requirements before configuring custom user registries:

- You must have the administrator or new admin role.
- Enable global security in your environment.
- Implement and build the `UserRegistry` interface and configure a custom registry.
- To configure custom user registries for multiple security domains, you must configure at least one security domain.

About this task

WebSphere Application Server security supports stand-alone custom registries in addition to the local operating system registry, standalone Lightweight Directory Access Protocol (LDAP) registries, and federated repositories for authentication and authorization. A stand-alone custom-implemented registry uses the UserRegistry Java interface as provided by the product. A stand-alone custom registry can support any type of account repository from a relational database, flat file, and so on. You can specify custom user registries at the global level and at the security domain.

When you configure a user registry in the global security configuration, the administrator does not specify a realm name for the user registry. The system determines the realm name from the security run time. The realm name for custom registries is set by the custom registry.

Use the following Jython command to make the user registry the active user registry in the global security configuration:

```
AdminTask.setAdminActiveSecuritySettings('-activeUserRegistry CustomUserRegistry')
```

Use the following Jython command to make the user registry the active user registry in the application security configuration:

```
AdminTask.setAppActiveSecuritySettings('-securityDomainName domain2 -activeUserRegistry CustomUserRegistry')
```

In security domains, you can configure a different realm for a user registry configuration. For example, you can configure two registries that use the same LDAP server listening on the same port, but use different base distinguished names (baseDN). This method supports the configuration to serve different sets of users and groups. To use this type of scenario, you must specify a realm name for each user registry configured for a domain. Multiple realms can exist in your configuration, and you can also specify a list of trusted realms. Communications between applications that use different realms is supported.

Use the following steps to configure custom user registries for your global security configuration and for multiple security domains:

- Configure custom user registries for global security configurations.
Use the `configureAdminCustomUserRegistry` command and the following optional parameters to configure a custom user registry in your global security configuration:

Parameter	Description	Data Type
-autoGenerateServerId	Specifies whether to automatically generate the server identity to use for internal process communication. To set a specific server identity, specify the -serverId parameter.	Boolean
-serverId	Specifies the user identity in the repository to use for internal process communication.	String
-serverIdPassword	Specifies the password that corresponds to the user identity.	String
-primaryAdminId	Specifies the name of the user with administrative privileges as defined in the registry. This parameter does not apply to security configurations. The user name must exist in the user registry repository.	String
-customRegClass	Specifies the class name that implements the UserRegistry interface in the <code>com.ibm.websphere.security</code> class.	String
-ignoreCase	Specifies whether to require case sensitive authorization. Specify true to ignore case during authorization.	Boolean
-customProperties	Specifies a list of attribute and value pairs to store as custom properties on the user registry object. Separate each attribute and value pair with a comma character (,) as the following syntax displays: "attribute1=value1","attribute2=value2"	String

Parameter	Description	Data Type
-verifyRegistry	Specifies whether to verify the user registry. The default value is true and verification is automatically performed.	Boolean

Use the following Jython example command to configure the local operating system registry for global security:

```
AdminTask.configureAdminCustomUserRegistry('-autoGenerateServerId true -primaryAdminId gsAdmin')
```

- Configure custom user registries for security domains.

1. Determine the name of the security domain to configure.

Use the listSecurityDomains command to list all security domains on the server, as the following Jython example demonstrates:

```
AdminTask.listSecurityDomains()
```

2. Configure a custom user registry for a security domain.

Use the configureAppCustomUserRegistry command and the following optional parameters to configure a local custom user registry:

Parameter	Description	Data type
-securityDomainName	Specifies the unique name that identifies the security domain of interest.	String
-realmName	Specifies the name of the realm of the user registry.	String
-customRegClass	Specifies the class name that implements the UserRegistry interface in the com.ibm.websphere.security class.	String
-ignoreCase	Specifies whether to require case sensitive authorization. Specify true to ignore case during authorization.	Boolean
-customProperties	Specifies a list of attribute and value pairs to store as custom properties on the user registry object. Separate each attribute and value pair with a comma character (,) as the following syntax displays: "attribute1=value1","attribute2=value2"	String
-verifyRegistry	Specifies whether to verify the user registry. The default value is true and verification is automatically performed.	Boolean

Use the following Jython example command to configure the local operating system user registry for the domain2 security domain:

```
AdminTask.configureAppCustomUserRegistry('-securityDomainName domain2 -realmName domain2Realm')
```

What to do next

Configuring JAAS login modules using scripting

Use this topic to use the wsadmin tool to configure and manage Java Authentication and Authorization Service (JAAS) login entries to allow communication between realms in a multiple security domain environment.

Before you begin

You must meet the following requirements before configuring local operating system user registries:

- You must have the administrator or new admin role.
 - Enable global security in your environment.
 - Configure multiple realms using security domains in your environment.
1. Launch the wsadmin scripting tool using the Jython scripting language.
 2. Configure a JAAS login module.

Use the `configureJAASLoginEntry` command to configure a Java Authentication and Authorization Service (JAAS) login entry in a security domain or in the global security configuration. You can use this command to modify existing JAAS login entries or to create new login entries.

Specify the following parameters to configure the JAAS login module:

Parameter	Description
<code>-loginEntryAlias</code>	Specifies an alias that identifies the JAAS login entry in the configuration. (String, required)
<code>-loginType</code>	Specifies the type of JAAS login entry of interest. Specify <code>system</code> for the system login type or <code>application</code> for the application login type. (String, required)
<code>-securityDomainName</code>	Specifies the name of the security configuration. If you do not specify a security domain name, the system updates the global security configuration. (String, optional)
<code>-loginModules</code>	Specifies a comma (,) separated list of login module class names. Specify the list in the order that the system calls them. (String, optional)
<code>-authStrategies</code>	<p>Optionally specifies the authentication behavior as authentication proceeds down the list of login modules. (String, optional)</p> <p>Specify one or many of the following values in a comma (,) separated list:</p> <ul style="list-style-type: none"> REQUIRED Specifies that the LoginModule module is required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list for each realm. REQUISITE Specifies that the LoginModule module is required to succeed. If authentication is successful, the process continues down the LoginModule list in the realm entry. If authentication fails, control immediately returns to the application. Authentication does not proceed down the LoginModule list. SUFFICIENT Specifies that the LoginModule module is not required to succeed. If authentication succeeds, control immediately returns to the application. Authentication does not proceed down the LoginModule list. If authentication fails, the process continues down the list. OPTIONAL Specifies that the LoginModule module is not required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list.

Use the `configureJAASLoginEntry` command to configure the JAAS login module, as the following Jython example demonstrates:

```
AdminTask.configureJAASLoginEntry(['[-securityDomainName testDomain
-loginType application -loginEntryAlias testLoginEntry -loginModules
"com.ibm.ws.security.common.auth.module.WSLoginModuleImpl" -authStrategies "REQUIRED"]'])
```

3. Set custom properties for the JAAS login module.

Use the `configureLoginModule` command to specify custom properties, modify the authentication strategy, or set the module to use a login module proxy. The following Jython command sets the debug and delegate custom properties for the `testLoginEntry` JAAS login entry:

```
AdminTask.configureLoginModule(['[-securityDomainName testDomain -loginType application
-loginEntryAlias testLoginEntry -loginModule com.ibm.ws.security.common.auth.module.WSLoginModuleImpl
-customProperties ["debug=true","delegate=WSLogin"]'])
```

4. Save your configuration changes.

Configuring Common Secure Interoperability authentication using scripting

Use this topic to use the `wsadmin` tool to configure inbound and outbound communications using the Common Secure Interoperability protocol. Common Secure Interoperability Version 2 (CSlv2) supports increased vendor interoperability and additional features.

Before you begin

You must meet the following requirements before configuring local operating system user registries:

- You must have the administrator or new admin role.
- Enable global security in your environment.
- Configure multiple realms using security domains in your environment.
- Configure CSI inbound communication authentication.

Inbound authentication refers to the configuration that determines the type of accepted authentication for inbound requests. This authentication is advertised in the interoperable object reference (IOR) that the client retrieves from the name server.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine the settings to specify for CSI inbound communication.

The `configureCSInbound` command configures various settings for CSI inbound communication. Review the following list of optional parameters to determine the attributes to set in your configuration:

Parameter	Description
<code>-securityDomainName</code>	Specifies the name of the security configuration. If you do not specify a security domain name, the command modifies the global security configuration. (String)
<code>-messageLevelAuth</code>	Specifies whether clients connecting to this server must specify a user ID and password. Specify <code>Never</code> to disable the user ID and password requirement. Specify <code>Supported</code> to accept a user ID and password. Specify <code>Required</code> to require a user ID and password. (String)
<code>-supportedAuthMechList</code>	Specifies the authentication mechanism to use. Specify <code>KRB5</code> for Kerberos authentication, <code>LTPA</code> for Lightweight Third-Party Authentication, <code>BasicAuth</code> for basic authentication, and <code>custom</code> to use your own authentication token implementation. You can specify more than one, separated by the pipe character (<code> </code>). (String)
<code>-clientCertAuth</code>	Specifies whether a client that connects to the server must connect using an SSL certificate. Specify <code>Never</code> to allow clients to connect without SSL certificates. Specify <code>Supported</code> to accept clients connecting with and without SSL certificates. Specify <code>Required</code> to require clients to use SSL certificate. (String)
<code>-transportLayer</code>	Specifies the transport layer support level. Specify <code>Never</code> to disable transport layer support. Specify <code>Supported</code> to enable transport layer support. Specify <code>Required</code> to require transport layer support. (String)
<code>-sslConfiguration</code>	Specifies the SSL configuration alias to use for inbound transport. (String)
<code>-enableIdentityAssertion</code>	Specifies whether to enable identity assertion. When using the identity assertion authentication method, the security token generated is a <code><wsse:UsernameToken></code> element that contains a <code><wsse:Username></code> element. Specify <code>true</code> for the <code>-enableIdentityAssertion</code> parameter to enable identity assertion. (Boolean)
<code>-trustedIdentities</code>	Specifies a list of trusted server identities, separated by the pipe character (<code> </code>). To specify a null value, set the value of the <code>-trustedIdentities</code> parameter as an empty string (<code>""</code>). (String)
<code>-statefulSession</code>	Specifies whether to enable a stateful session. Specify <code>true</code> to enable a stateful session. (Boolean)
<code>-enableAttributePropagation</code>	Specifies whether to enable security attribute propagation. Security attribute propagation allows the application server to transport authenticated subject contents and security context information from one server to another in your configuration. Specify <code>true</code> to enable security attribute propagation. (Boolean)

3. Configure CSI inbound communication authentication.

The `configureCSInbound` command configures the CSIv2 Inbound authentication on a security domain or on the global security configuration. When configuring CSI Inbound in a security domain for the first time, the CSI objects are copied from global security. Then, the changes are applied to configuration.

Use the `configureCSInbound` command to configure CSI inbound authentication for a security domain or the global security configuration, as the following Jython example demonstrates:

```
AdminTask.configureCSInbound('-securityDomainName testDomain -messageLevelAuth Supported  
-supportedAuthMechList KRB5|LTPA -clientCertAuth Supported -statefulSession true')
```

4. Save your configuration changes.

- Configure CSI outbound communication authentication.

Outbound authentication refers to the configuration that determines the type of authentication that is performed for outbound requests to downstream servers.

1. Launch the wsadmin scripting tool using the Jython scripting language.

2. Determine the settings to specify for CSI outbound communication.

The `configureCSIOutbound` command configures various settings for CSI outbound communication. Review the following list of optional parameters to determine the attributes to set in your configuration:

Parameter	Description
<code>-securityDomainName</code>	Specifies the name of the security configuration. If you do not specify a security domain name, the command modifies the global security configuration. (String)
<code>-enableAttributePropagation</code>	Specifies whether to enable security attribute propagation. Security attribute propagation allows the application server to transport authenticated subject contents and security context information from one server to another in your configuration. Specify <code>true</code> to enable security attribute propagation. (Boolean)
<code>-enableIdentityAssertion</code>	Specifies whether to enable identity assertion. When using the identity assertion authentication method, the security token generated is a <code><wsse:UsernameToken></code> element that contains a <code><wsse:Username></code> element. Specify <code>true</code> for the <code>-enableIdentityAssertion</code> parameter to enable identity assertion. (Boolean)
<code>-useServerIdentity</code>	Specifies whether to use the server identity to establish trust with the target server. Specify <code>true</code> to use the server identity. (Boolean)
<code>-trustedId</code>	Specifies the trusted identity that the application server uses to establish trust with the target server. (String)
<code>-trustedIdentityPassword</code>	Specifies the password of the trusted server identity. (String)
<code>-messageLevelAuth</code>	Specifies whether clients connecting to this server must specify a user ID and password. Specify <code>Never</code> to disable the user ID and password requirement. Specify <code>Supported</code> to accept a user ID and password. Specify <code>Required</code> to require a user ID and password. (String)
<code>-supportedAuthMechList</code>	Specifies the authentication mechanism to use. Specify <code>KRB5</code> for Kerberos authentication, <code>LTPA</code> for Lightweight Third-Party Authentication, <code>BasicAuth</code> for basic authentication, and <code>custom</code> to use your own authentication token implementation. You can specify more than one, separated by the pipe character (<code> </code>). (String)
<code>-clientCertAuth</code>	Specifies whether a client that connects to the server must connect using an SSL certificate. Specify <code>Never</code> to allow clients to connect without SSL certificates. Specify <code>Supported</code> to accept clients connecting with and without SSL certificates. Specify <code>Required</code> to require clients to use SSL certificate. (String)
<code>-transportLayer</code>	Specifies the transport layer support level. Specify <code>Never</code> to disable transport layer support. Specify <code>Supported</code> to enable transport layer support. Specify <code>Required</code> to require transport layer support. (String)
<code>-sslConfiguration</code>	Specifies the SSL configuration alias to use for inbound transport. (String)
<code>-statefulSession</code>	Specifies whether to enable a stateful session. Specify <code>true</code> to enable a stateful session. (Boolean)
<code>-enableOutboundMapping</code>	Specifies whether to enable custom outbound identity mapping. Specify <code>true</code> to enable custom outbound identity mapping. (Boolean)
<code>-trustedTargetRealms</code>	Specifies a list of target realms to trust. Separate each realm name with the pipe character (<code> </code>). (String)

3. Configure CSI outbound communication authentication.

The `configureCSIOutbound` command configures the CSIv2 outbound authentication in a security domain or in the global security configuration. When configuring CSI outbound authentication in a security domain for the first time, the application server copies the CSI objects from global security. Then, the application server applies the changes to that configuration.

Use the `configureCSIOutbound` command to configure CSI outbound authentication for a security domain or the global security configuration, as the following Jython example demonstrates:

```
AdminTask.configureCSIOutbound('-securityDomainName testDomain -enableIdentityAssertion true
-trustedId myID -trustedIdentityPassword myPassword123 -messageLevelAuth Required
-trustedTargetRealms realm1|realm2|realm3')
```

4. Save your configuration changes.

Configuring trust association using scripting

Use the `wsadmin` tool to configure and manage trust association configurations in a multiple security domain environment. Trust association enables the integration of the application server security and third-party security servers. More specifically, a reverse proxy server can act as a front-end authentication server while the product applies its own authorization policy onto the resulting credentials that are passed by the proxy server.

Before you begin

You must meet the following requirements before configuring a trust association:

- You must have the administrator or new admin role.
 - Enable global security in your environment.
 - Configure multiple realms using security domains in your environment.
1. Launch the wsadmin scripting tool using the Jython scripting language.
 2. Configure a trust association.

Use the `configureTrustAssociation` command to enable the trust association. You can also use this command to create or modify a trust association interceptor.

The following Jython command creates a trust association for the `testDomain` security domain and configures the trust association to act as a reverse proxy server:

```
AdminTask.configureTrustAssociation('-securityDomainName testDomain -enable true')
```

3. Configure the trust association interceptor.

Use the `configureInterceptor` command to modify an existing interceptor. The following Jython command uses a WebSEAL interceptor to configure single sign-on for the `testDomain` security domain:

```
AdminTask.configureInterceptor('[-interceptor com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus  
-securityDomainName testDomain -customProperties ["com.ibm.websphere.security.trustassociation.types=webseal",  
"com.ibm.websphere.security.webseal.loginId=websealLoginID","com.ibm.websphere.security.webseal.id=iv-user"]]')
```

4. Save your configuration changes.

Mapping resources to security domains using scripting

Use this topic to assign management resources to security domains. Set management resources to your security domains to customize your security configuration for a cell, server, or cluster.

Before you begin

Users assigned to the administrator role can configure security domains. Verify that you have the appropriate administrative role before configuring security domains. Also, create a security domain, or copy an existing security domain before assigning resources to a security domain.

About this task

After creating a security domain, you can map management resources to the security domain. You can assign resources to a security domain at the server, cell, and cluster level. Use the following steps to assign a resource to a security domain:

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine which security domain to map a resource.

Use the `listSecurityDomains` command to view a list of security domains in your configuration. Specify `true` for the optional `-listDescription` parameter to list the description for each security domain, as the following Jython example demonstrates:

```
print AdminTask.listSecurityDomains('-listDescription true')
```

The command returns the following example attribute list output:

```
{{name myDomain}  
{description {security domain for administrative applications}}}  
{{name domain2}  
{description {new domain for cell123}}}
```

3. Assign a resource to a security domain.

Use the `mapResourceToSecurityDomain` command to assign a management resource to the security domain. For example, use the following Jython command to secure all applications on the `server1` cell with the security attributes in the `domain2` security domain:

```
AdminTask.mapResourceToSecurityDomain('-securityDomainName domain2 -resourceName Cell=myCell:Node=myNode:Server=server1')
```

4. Save your configuration changes.

Results

Your security domain is updated in your configuration. All applications in the specified resource use the security attributes specified by the security domain. If the security domain does not contain all security attributes, then the missing attributes are obtained from the global security configuration.

What to do next

Restart each resource that you assigned to a security domain.

Removing resources from security domains using scripting

Use this topic to remove management resources from security domains. Remove all resources from a security domain before deleting the security domain from your configuration.

Before you begin

Users assigned to the administrator role can configure security domains. Verify that you have the appropriate administrative role before configuring security domains.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine the security domain to edit.

Use the `listSecurityDomains` command to view a list of security domains in your configuration. Specify `true` for the optional `-listDescription` parameter to list the description for each security domain, as the following Jython example demonstrates:

```
print AdminTask.listSecurityDomains('-listDescription true')
```

The command returns the following example output:

```
myDomain - security domain for administrative applications  
domain2 - new domain for cell123
```

3. Verify the management resources that are assigned to the security domain.

Use the `listResourcesInSecurityDomain` command to view a list of resources that are mapped to the security domain of interest, as the following Jython example demonstrates:

```
print AdminTask.listResourcesInSecurityDomain('-securityDomainName domain2')
```

The command returns the following example output:

```
"Cell=myhostCell101"
```

4. Remove the resource from the security domain.

Use the `removeResourceFromSecurityDomain` command to remove a management resource from the security domain. For example, use the following Jython command to remove the `Cell101` cell resource from the `domain2` security domain:

```
AdminTask.removeResourceFromSecurityDomain('-securityDomainName domain2 -resourceName Cell=myhostCell101')
```

5. Save your configuration changes.

What to do next

Restart each management resource that you removed from a security domain.

Removing security domains using scripting

Use this topic to delete security domains from your configuration using the wsadmin tool. Remove security domains that are not needed in your security configuration.

Before you begin

Users assigned to the administrator role can configure security domains. Verify that you have the appropriate administrative role before configuring security domains. A security domain must exist in your configuration.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine the security domain to delete.

Use the `listSecurityDomains` to view a list of security domains in your configuration. Specify `true` for the optional `-listDescription` parameter to list the description for each security domain, as the following Jython example demonstrates:

```
print AdminTask.listSecurityDomains('-listDescription true')
```

The command returns the following example output:

```
{{name myDomain}
 {description {security domain for administrative applications}}}
 {{name domain2}
 {description {new domain for cell123}}}
```

3. Verify that no resources are assigned to the security domain to delete.

You can use this step to manually remove resources from the security domain of interest. You do not need to complete this step if you want to delete the security domain and each assigned resource. Use the `listResourcesInSecurityDomain` command to view a list of resources that are mapped to the security domain of interest, as the following Jython example demonstrates:

```
print AdminTask.listResourcesInSecurityDomain('-securityDomainName domain2')
```

If the command returns the name of a resource, use the `removeResourceFromSecurityDomain` command to remove a resource from the security domain. For example, use the following Jython command to remove the `Cell101` cell resource from the `domain2` security domain:

```
"AdminTask.removeResourceFromSecurityDomain('-securityDomainName domain2 -resourceName Cell=myhostCell101')"
```

4. Delete the security domain from your configuration.

Use the `deleteSecurityDomain` command to delete the security domain. If a resource associated with the domain was deleted from the system but the mapping was not removed from the domain, specify the optional `-force` parameter to remove the domain, as the following Jython example demonstrates:

```
AdminTask.deleteSecurityDomain('-securityDomainName domain2 -force true')
```

5. Save your configuration changes.

Removing user registries using scripting

You can use the wsadmin tool to remove user registries from global security or security domain configurations. Use the steps in this topic to remove Lightweight Directory Access Protocol (LDAP), local operating system, custom, or federated repository user registries from your global security or security domain configurations.

Before you begin

You must meet the following requirements before configuring local operating system user registries:

- You must have the administrator or new admin role.
- Enable global security in your environment.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine the registry to remove.

Use the `getUserRegistryInfo` command to display information about a user registry from the global security configuration or in a security domain. You must specify the type of user registry of interest. Valid values are `LDAPUserRegistry`, `WIMUserRegistry`, `CustomUserRegistry`, and `LocalOSUserRegistry`.

The following command returns a list of values in the local operating system user registry object for the `domain2` security domain, as the following example Jython demonstrates:

```
AdminTask.getUserRegistryInfo('-securityDomainName domain2 -userRegistryType LocalOSUserRegistry')
```

3. Determine whether the registry of interest is the active user registry.

You cannot remove the active user registry. Use the `getActiveSecuritySettings` command to see check if the user registry is the active user registry before removing it.

4. Remove the registry of interest.

Use the `unconfigureUserRegistry` command to remove the registry of interest. If you remove the user registry from the global security configuration, then the command reduces the registry object to the minimum values for the configuration. If you remove the user registry from a security domain, then the command removes the configuration object from the security domain. The following Jython example removes the local operating system user registry configuration from the `domain2` security domain:

```
AdminTask.unconfigureUserRegistry('-securityDomainName domain2 -userRegistryType LocalOSUserRegistry')
```

5. Save your configuration changes.

SecurityDomainCommands command group for the AdminTask object

You can use the Jython scripting language to configure and administer security domains with the `wsadmin` tool. Use the commands and parameters in the `SecurityDomainCommands` group to create and manage security domains, assign servers and clusters to security domains as resources, and to query the security domain configuration.

Use the following commands to administer the security domain configuration:

- `copySecurityDomain`
- `copySecurityDomainFromGlobalSecurity`
- `createSecurityDomain`
- `deleteSecurityDomain`
- `getSecurityDomainForResource`
- `listResourcesInSecurityDomain`
- `listSecurityDomains`
- “`listSecurityDomainsForResources`” on page 897
- `mapResourceToSecurityDomain`
- `modifySecurityDomain`
- `removeResourceFromSecurityDomain`

copySecurityDomain

The `copySecurityDomain` command creates a new security domain by copying an existing security domain. If the security configuration defines an active user registry, provide a realm name for the newly create security domain. If you do not specify a realm name, the system creates a realm name.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the new security domain that the system creates by copying another security domain. (String)

-copyFromSecurityDomainName

Specifies the name of the existing security domain that the system uses to create the new security domain. (String)

Optional parameters

-securityDomainDescription

Specifies a description for the new security domain. (String)

-realmName

Specifies the name of the realm in the new security domain. The system creates a name for the realm if you do not specify a value for this parameter. (String)

Return value

The command returns the configuration ID of the new security domain.

Batch mode example usage

- Using Jython string:

```
AdminTask.copySecurityDomain('-securityDomainName copyOfDomain2 -copyFromSecurityDomainName Domain2')
```

- Using Jython list:

```
AdminTask.copySecurityDomain('-securityDomainName', 'copyOfDomain2', '-copyFromSecurityDomainName', 'Domain2')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.copySecurityDomain('-interactive')
```

copySecurityDomainFromGlobalSecurity

The `copySecurityDomainFromGlobalSecurity` command creates a security domain by copying the global security configuration. If an active user registry exists for the global security configuration, provide a realm name for the newly created security domain. If you do not specify a realm name, then the system creates a realm name.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the new security domain that the system copies from the global security configuration. (String)

Optional parameters

-securityDomainDescription

Specifies a description for the new security domain. (String)

-realmName

Specifies the name of the realm in the new security configuration. The system creates a name for the realm if you do not specify a value for the `-realmName` parameter. (String)

Return value

The command returns the configuration ID of the new security domain.

Batch mode example usage

- Using Jython string:

```
AdminTask.copySecurityDomainFromGlobalSecurity('-securityDomainName GSCopy -securityDomainDescription "copy of global security" -realmName myRealm')
```

- Using Jython list:

```
AdminTask.copySecurityDomainFromGlobalSecurity('-securityDomainName', 'GSCopy', '-securityDomainDescription', '"copy of global security"', '-realmName myRealm')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.copySecurityDomainFromGlobalSecurity('-interactive')
```

createSecurityDomain

The createSecurityDomain command creates the security domain-security.xml and domain-security-map.xml files under the profile_root/config/cells/cellName/securityDomain/configurationName directory. The system creates an empty domain-security.xml file.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the new security domain to create. (String)

Optional parameters

-securityDomainDescription

Specifies a description of the new security domain. (String)

Return value

The command returns the configuration ID of the new security domain.

Batch mode example usage

- Using Jython string:

```
AdminTask.createSecurityDomain('-securityDomainName newDomain -securityDomainDescription "new security domain"')
```

- Using Jython list:

```
AdminTask.createSecurityDomain('-securityDomainName', 'newDomain', '-securityDomainDescription',  
"new security domain")
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createSecurityDomain('-interactive')
```

deleteSecurityDomain

The deleteSecurityDomain command removes the domain-security.xml and domain-security-map.xml files from the security domain directory. The command returns an error if resources are mapped to the security domain of interest. To delete the security domain when resources are mapped to the security domain of interest, specify the value for the -force parameter as true.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security domain to delete. (String)

Optional parameters

-force

Specifies that the system deletes the security domain without checking for resources that are

associated with the domain. Use this option when the resources in the security domains are not valid resources. The default value for the `-force` parameter is `false`. (Boolean)

Return value

The command does not return output if the system successfully removes the security domain configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteSecurityDomain('-securityDomainName mySecurityDomain -force true')
```

- Using Jython list:

```
AdminTask.deleteSecurityDomain('-securityDomainName', 'mySecurityDomain', '-force', 'true')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteSecurityDomain('-interactive')
```

getSecurityDomainForResource

The `getSecurityDomainForResource` command displays the security domain for a specific resource. If the resource is not mapped to a domain, the command does not return output.

Target object

None.

Required parameters

-resourceName

Specifies the name of the resource of interest. Specify the value in the following format:
`Cell=:Node=myNode:Server=myServer` (String)

Optional parameters

-getEffectiveDomain

Specifies whether the command returns the effective domain of the resource if the resource is not directly mapped to a domain. The default value is `true`. Specify `false` if you do not want to display the effective domain if the resource is not directly mapped to a domain. (Boolean)

Return value

The command returns the security domain name as a string.

Batch mode example usage

- Using Jython string:

```
AdminTask.getSecurityDomainForResource('-resourceName Cell=:Node=myNode:Server=myServer')
```

- Using Jython list:

```
AdminTask.getSecurityDomainForResource('-resourceName', 'Cell=:Node=myNode:Server=myServer')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getSecurityDomainForResource('-interactive')
```

listResourcesInSecurityDomain

The `listResourcesInSecurityDomain` command displays the servers or clusters that are associated with a specific security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security domain of interest. (String)

-expandCell

Specifies whether to display the servers in the cell. Specify `true` to display the specific servers, or specify `false` to list the cell information only. (Boolean)

Return value

The command returns an array that contains the names of the resources that are mapped to the security domain of interest in the format: `Cell=<cell name>;Node=<node name>;Server=<server name>`.

Batch mode example usage

- Using Jython string:

```
AdminTask.listResourcesInSecurityDomain('-securityDomainName myDomain')
```

- Using Jython list:

```
AdminTask.listResourcesInSecurityDomain('-securityDomainName', 'myDomain')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listResourcesInSecurityDomain('-interactive')
```

listSecurityDomains

The `listSecurityDomains` command lists each security domain configured for the server.

Target object

None.

Optional parameters

-listDescription

Specifies whether to display the description of the security domains. Specify `true` to display the descriptions of the security domains. (Boolean)

-doNotDisplaySpecialDomains

Specifies whether to exclude special domains. Specify `true` to exclude the special domains in the command output, or `false` to display the special domains. (Boolean)

Return value

The command returns an array that contains the names of security domains that are configured for the server. The command returns an array of attribute lists that contain the name and description for each security domain if the `-listDescription` parameter is specified.

Batch mode example usage

- Using Jython string:

```
AdminTask.listSecurityDomains('-listDescription true')
```

- Using Jython list:

```
AdminTask.listSecurityDomains('-listDescription', 'true')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listSecurityDomains('-interactive')
```

listSecurityDomainsForResources

The `listSecurityDomainsForResources` command lists the security domains that are associated with the resources of interest.

Target object

None.

Required parameters

-resourceNames

Specifies one or more resources for which the command returns the associated security domains. Specify each resource separated by the plus sign character (+). (String)

Return value

The command returns the list of resources specified by the `-resourceNames` parameter and the security domains to which each resource is mapped.

Batch mode example usage

- Using Jython string:

```
AdminTask.listSecurityDomainsForResources('-resourceNames resource1+resource2+resource3')
```

- Using Jython list:

```
AdminTask.listSecurityDomainsForResources('-resourceNames', 'resource1+resource2+resource3')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listSecurityDomainsForResources('-interactive')
```

mapResourceToSecurityDomain

The `mapResourceToSecurityDomain` command maps a resource to a security domain. The system adds an entry for each resource to the `domain-security-map.xml` file.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security domain of interest. (String)

-resourceName

Specifies the name of the resource to which the system maps the security domain of interest. Specify the value in the following format: `Cell=:Node=myNode:Server=myServer` (String)

Return value

The command does not return output if the system successfully assigns the resource to the security domain of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.mapResourceToSecurityDomain('-securityDomainName mySecurityDomain -resourceName  
-resourceName Cell=:Node=myNode:Server=myServer')
```

- Using Jython list:

```
AdminTask.mapResourceToSecurityDomain('-securityDomainName', 'mySecurityDomain', '-resourceName',  
'-resourceName Cell=:Node=myNode:Server=myServer')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.mapResourceToSecurityDomain('-interactive')
```

modifySecurityDomain

The `modifySecurityDomain` command changes the description of a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security domain to edit. (String)

Optional parameters

-securityDomainDescription

Specifies the new description for the security domain of interest. (String)

Return value

The command does not return output if the system successfully modifies the security domain.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifySecurityDomain('-securityDomainName myDomain -securityDomainDescription  
"my new description"')
```

- Using Jython list:

```
AdminTask.modifySecurityDomain('-securityDomainName', 'myDomain', '-securityDomainDescription',=  
"my new description"')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifySecurityDomain('-interactive')
```

removeResourceFromSecurityDomain

The `removeResourceFromSecurityDomain` command removes a resource from a security domain mapping. The command removes the resource entry from the `domain-security-map.xml` file.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security domain from which to remove the resource. (String)

-resourceName

Specifies the name of the resource to remove. Specify the value in the following format:
`Cell=:Node=myNode:Server=myServer` (String)

Return value

The command does not return output if the system successfully removes the resource from the security domain.

Batch mode example usage

- Using Jython string:

```
AdminTask.removeResourceFromSecurityDomain('-securityDomainName myDomain -resourceName  
Cell=:Node=myNode:Server=myServer')
```

- Using Jython list:

```
AdminTask.removeResourceFromSecurityDomain('-securityDomainName', 'myDomain', '-resourceName',  
'Cell=:Node=myNode:Server=myServer')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.removeResourceFromSecurityDomain('-interactive')
```

SecurityConfigurationCommands command group for the AdminTask object

You can use the Jython scripting language to configure security with the `wsadmin` tool. Use the commands and parameters in the `SecurityConfigurationCommands` group to configure and manage user registries, single sign-on, data entries, trust association, login modules, and interceptors.

Use the following command to administer user registry configurations:

- `configureAdminCustomUserRegistry`
- `configureAdminLDAPUserRegistry`
- `configureAdminLocalOSUserRegistry`
- `configureAdminWIMUserRegistry`
- `configureAppCustomUserRegistry`
- `configureAppLDAPUserRegistry`
- `configureAppLocalOSUserRegistry`
- `configureAppWIMUserRegistry`
- `getLTPATimeout`
- `setLTPATimeout`
- `getUserRegistryInfo`

- `unconfigureUserRegistry`

Use the following commands to administer Java Authentication and Authorization Service (JAAS) login configurations:

- `configureLoginEntry`
- `configureLoginModule`
- `getJAASLoginEntryInfo`
- `listJAASLoginEntries`
- `listLoginModules`
- `unconfigureJAASLoginEntry`
- `unconfigureLoginModule`

Use the following commands to administer data entry configurations:

- `createAuthDataEntry`
- `deleteAuthDataEntry`
- `getAuthDataEntry`
- `listAuthDataEntries`
- `modifyAuthDataEntry`

Use the following commands to administer Common Secure Interoperability Version 2 (CSIv2) configurations:

- `configureCSInbound`
- `configureCSIOutbound`
- `getCSInboundInfo`
- `getCSIOutboundInfo`
- `unconfigureCSInbound`
- `unconfigureCSIOutbound`

Use the following commands to administer trust association configurations:

- `configureInterceptor`
- `configureTrustAssociation`
- `getTrustAssociationInfo`
- `listInterceptors`
- `unconfigureInterceptor`
- `unconfigureTrustAssociation`

Use the following commands to manage your security configuration:

- `configureAuthzConfig`
- `configureSingleSignon`
- `getActiveSecuritySettings`
- “`getAuthzConfigInfo`” on page 934
- “`getSingleSignon`” on page 935
- `setAdminActiveSecuritySettings`
- `setAppActiveSecuritySettings`
- `unconfigureAuthzConfig`
- `unsetAppActiveSecuritySettings`

configureAdminCustomUserRegistry

The `configureAdminCustomUserRegistry` command configures a custom user registry in the global security configuration.

Target object

None.

Optional parameters

-autoGenerateServerId

Specifies whether the command automatically generates the server identity that the system uses for internal process communication. Specify `true` to automatically generate the server identity. (Boolean)

-serverId

Specifies the server identity in the repository that the system uses for internal process communication. (String)

-serverIdPassword

Specifies the password that corresponds to the server identity. (String)

-primaryAdminId

Specifies the name of the user with administrative privileges that is defined in the registry. This parameter does not apply to security configurations. (String)

-customRegClass

Specifies the class name that implements the `UserRegistry` interface in `com.ibm.websphere.security` property. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAdminCustomUserRegistry('-autoGenerateServerId true -serverIdPassword password4server -primaryAdminId serverAdmin')
```

- Using Jython list:

```
AdminTask.configureAdminCustomUserRegistry(['-autoGenerateServerId', 'true', '-serverIdPassword', 'password4server', '-primaryAdminId', 'serverAdmin'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAdminCustomUserRegistry('-interactive')
```

configureAdminLDAPUserRegistry

The configureAdminLDAPUserRegistry command configures a Lightweight Directory Access Protocol (LDAP) user registry in the global security configuration.

Target object

None.

Optional parameters

-autoGenerateServerId

Specifies whether the command automatically generates the server identity used for internal process communication. Specify true to automatically generate the server identity. (Boolean)

-serverId

Specifies the server identity in the repository that the system uses for internal process communication. (String)

-serverIdPassword

Specifies the password that corresponds to the server identity. (String)

-primaryAdminId

Specifies the name of the user with administrative privileges that is defined in the registry. This parameter does not apply to security configurations. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to true, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to false to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-ldapServerType

Specifies the type of LDAP server. The default type is IBM_DIRECTORY_SERVER. (String)

Specify one of the following valid values:

- IBM_DIRECTORY_SERVER
- IPLANET
- NETSCAPE
- NDS
- DOMIN0502
- SECUREWAY
- ACTIVE_DIRECTORY
- CUSTOM

-ldapHost

Specifies the host name of the LDAP server. (String)

-ldapPort

Specifies the port that the system uses to access the LDAP server. The default value is 389. (String)

-baseDN

Specifies the base distinguished name (DN) of the directory service, which indicates the starting point for LDAP searches of the directory service. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed. (String)

-bindDN

Specifies the distinguished name for the application server, which is used to bind to the directory service. (String)

-bindPassword

Specifies the binding DN password for the LDAP server. (String)

-searchTimeout

Specifies the timeout value in seconds for an LDAP server to respond before stopping a request. The default value is 120 seconds. (Long)

-reuseConnection

Specifies whether the server reuses the LDAP connection. By default, this option is enabled. Specify `false` for this parameter only in rare situations where a router is used to distribute requests to multiple LDAP servers and when the router does not support affinity. (Boolean)

Note: When you disable the reuse of the LDAP connection, the application server creates a new LDAP connection for every LDAP search request. This situation impacts system performance if your environment requires extensive LDAP calls. This option is provided because the router is not sending the request to the same LDAP server. The option is also used when the idle connection timeout value or firewall timeout value between the application server and LDAP is too small.

-userFilter

Specifies the LDAP filter clause that the system uses to search the user registry for users. The default value is the default user filter for the LDAP server type. (String)

-groupFilter

Specifies the LDAP filter clause that the system uses to search the user registry for groups. The default value is the default group filter for the LDAP server type. (String)

-userIdMap

Specifies the LDAP filter that maps the short name of a user to an LDAP entry. The default value is the default user filter for the LDAP server type. (String)

-groupIdMap

Specifies the LDAP filter that maps the short name of a group to an LDAP entry. The default value is the default group filter for the LDAP server type. (String)

-groupMemberIdMap

Specifies the LDAP filter that identifies users to group memberships. (String)

-certificateMapMode

Specifies whether to map X.509 certificates into an LDAP directory by `EXACT_DN` or `CERTIFICATE_FILTER`. Specify `CERTIFICATE_FILTER` to use the specified certificate filter for the mapping. (String)

-certificateFilter

Specifies the filter certificate mapping property for the LDAP filter. The filter is used to map attributes in the client certificate to entries in the LDAP registry. (String)

The syntax or structure of this filter is: `(&(uid=${SubjectCN})(objectclass=inetOrgPerson))`. The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. The right side must begin with a dollar sign (\$) and open bracket ({} and end with a close bracket (}). You can use the following certificate attribute values on the right side of the filter specification. The case of the strings is important:

- `${UniqueKey}`
- `${PublicKey}`
- `${Issuer}`

- `${NotAfter}`
- `${NotBefore}`
- `${SerialNumber}`
- `${SigAlgName}`
- `${SigAlgOID}`
- `${SigAlgParams}`
- `${SubjectCN}`
- `${Version}`

-krbUserFilter

Specifies the default value is the default user filter for the LDAP server type. (String)

-nestedGroupSearch

Specifies whether to perform a recursive nested group search. Specify `true` to perform a recursive nested group search, or specify `false` to disable recursive nested group searching. (Boolean)

-sslEnabled

Specifies whether to enable Secure Sockets Layer (SSL). Specify `true` to enable an SSL connection to the LDAP server. (Boolean)

-sslConfig

Specifies the SSL configuration alias to use for the secure LDAP connection. (String)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAdminCustomUserRegistry('-autoGenerateServerId true -serverIdPassword password4server -primaryAdminId serverAdmin -ldapServerType NETSCAPE -ldapHost 195.168.1.1')
```

- Using Jython list:

```
AdminTask.configureAdminCustomUserRegistry(['-autoGenerateServerId', 'true', '-serverIdPassword', 'password4server', '-primaryAdminId', 'serverAdmin', '-ldapServerType', 'NETSCAPE', '-ldapHost', '195.168.1.1'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAdminLDAPUserRegistry('-interactive')
```

configureAdminLocalOSUserRegistry

The `configureAdminLocalOSUserRegistry` command configures a local operating system user registry in the global security configuration.

Target object

None.

Optional parameters

-autoGenerateServerId

Specifies whether the command automatically generates the server identity used for internal process communication. Specify `true` to automatically generate the server identity. (Boolean)

-serverId

Specifies the server identity in the repository that the system uses for internal process communication. (String)

-serverIdPassword

Specifies the password that corresponds to the server identity. (String)

-primaryAdminId

Specifies the name of the user with administrative privileges that is defined in the registry. This parameter does not apply to security configurations. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAdminLocalOSUserRegistry('-autoGenerateServerId true -serverIdPassword password4server  
-primaryAdminId serverAdmin')
```

- Using Jython list:

```
AdminTask.configureAdminLocalOSUserRegistry(['autoGenerateServerId', 'true', '-serverIdPassword', 'password4server',  
'-primaryAdminId', 'serverAdmin'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAdminLocalOSUserRegistry('-interactive')
```

configureAdminWIMUserRegistry

The `configureAdminWIMUserRegistry` command configures a federated repository user registry in the administrative security configuration.

Target object

None.

Optional parameters

-autoGenerateServerId

Specifies whether the command automatically generates the server identity used for internal process communication. Specify `true` to automatically generate the server identity. (Boolean)

-serverId

Specifies the server identity in the repository that the system uses for internal process communication. (String)

-serverIdPassword

Specifies the password that corresponds to the server identity. (String)

-primaryAdminId

Specifies the name of the user with administrative privileges that is defined in the registry. This parameter does not apply to security configurations. (String)

-realmName

Specifies the realm of the user registry. The system automatically generates a realm name if you do not specify a value for the `-realmName` parameter. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAdminWIMUserRegistry('-autoGenerateServerId true -serverIdPassword password4server
-primaryAdminId serverAdmin')
```

- Using Jython list:

```
AdminTask.configureAdminWIMUserRegistry(['autoGenerateServerId', 'true', '-serverIdPassword', 'password4server',
'-primaryAdminId', 'serverAdmin'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAdminWIMUserRegistry('-interactive')
```

configureAppCustomUserRegistry

The `configureAppCustomUserRegistry` command configures a custom user registry in an application security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Optional parameters

-realmName

Specifies the realm of the user registry. The system automatically generates a realm name if you do not specify a value for the `-realmName` parameter. (String)

-customRegClass

Specifies the class name that implements the `UserRegistry` interface in `com.ibm.websphere.security` property. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAppCustomUserRegistry('-securityDomainName testDomain -realmName server_name.domain:port_number')
```

- Using Jython list:

```
AdminTask.configureAppCustomUserRegistry(['-securityDomainName', 'testDomain', '-realmName', 'server_name.domain:port_number'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAppCustomUserRegistry('-interactive')
```

configureAppLDAPUserRegistry

The `configureAppLDAPUserRegistry` command configures LDAP user registries in a security configuration or a global security configuration.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Optional parameters

-realmName

Specifies the realm of the user registry. The system automatically generates a realm name if you do not specify a value for the `-realmName` parameter. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID.

If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-ldapServerType

Specifies the type of LDAP server. The default type is `IBM_DIRECTORY_SERVER`. (String)

Specify one of the following valid values:

- `IBM_DIRECTORY_SERVER`
- `IPLANET`
- `NETSCAPE`
- `NDS`
- `DOMINO502`
- `SECUREWAY`
- `ACTIVE_DIRECTORY`
- `CUSTOM`

-ldapHost

Specifies the host name of the LDAP server. (String)

-ldapPort

Specifies the port that the system uses to access the LDAP server. The default value is 389. (String)

-baseDN

Specifies the base distinguished name (DN) of the directory service, which indicates the starting point for LDAP searches of the directory service. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed. (String)

-bindDN

Specifies the distinguished name for the application server, which is used to bind to the directory service. (String)

-bindPassword

Specifies the binding DN password for the LDAP server. (String)

-searchTimeout

Specifies the timeout value in seconds for an LDAP server to respond before stopping a request. The default value is 120 seconds. (Long Integer)

-reuseConnection

Specifies whether the server reuses the LDAP connection. By default, this option is enabled. Specify `false` for this parameter only in rare situations where a router is used to distribute requests to multiple LDAP servers and when the router does not support affinity. (Boolean)

Note: When you disable the reuse of the LDAP connection, the application server creates a new LDAP connection for every LDAP search request. This situation impacts system performance if your environment requires extensive LDAP calls. This option is provided because the router is not sending the request to the same LDAP server. The option is also used when the idle connection timeout value or firewall timeout value between the application server and LDAP is too small.

-userFilter

Specifies the LDAP filter clause that the system uses to search the user registry for users. The default value is the default user filter for the LDAP server type. (String)

-groupFilter

Specifies the LDAP filter clause that the system uses to search the user registry for groups. The default value is the default group filter for the LDAP server type. (String)

-userIdMap

Specifies the LDAP filter that maps the short name of a user to an LDAP entry. The default value is the default user filter for the LDAP server type. (String)

-groupIdMap

Specifies the LDAP filter that maps the short name of a group to an LDAP entry. The default value is the default group filter for the LDAP server type. (String)

-groupMemberIdMap

Specifies the LDAP filter that identifies users to group memberships. (String)

-certificateMapMode

Specifies whether to map X.509 certificates into an LDAP directory by EXACT_DN or CERTIFICATE_FILTER. Specify CERTIFICATE_FILTER to use the specified certificate filter for the mapping. (String)

-certificateFilter

Specifies the filter certificate mapping property for the LDAP filter. The filter is used to map attributes in the client certificate to entries in the LDAP registry. (String)

The syntax or structure of this filter is: (&(uid=\${SubjectCN})(objectclass=inetOrgPerson)). The left side of the filter specification is an LDAP attribute that depends on the schema that your LDAP server is configured to use. The right side of the filter specification is one of the public attributes in your client certificate. The right side must begin with a dollar sign (\$) and open bracket (()) and end with a close bracket (}). You can use the following certificate attribute values on the right side of the filter specification. The case of the strings is important:

- \${UniqueKey}
- \${PublicKey}
- \${Issuer}
- \${NotAfter}
- \${NotBefore}
- \${SerialNumber}
- \${SigAlgName}
- \${SigAlgOID}
- \${SigAlgParams}
- \${SubjectCN}
- \${Version}

-krbUserFilter

Specifies the default value is the default user filter for the LDAP server type. (String)

-nestedGroupSearch

Specifies whether to perform a recursive nested group search. Specify true to perform a recursive nested group search, or specify false to disable recursive nested group searching. (Boolean)

-sslEnabled

Specifies whether to enable Secure Sockets Layer (SSL). Specify true to enable an SSL connection to the LDAP server. (Boolean)

-sslConfig

Specifies the SSL configuration alias to use for the secure LDAP connection. (String)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: "attr1=value1","attr2=value2" (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAppLDAPUserRegistry('-securityDomainName testDomain -ldapServerType NETSCAPE -ldapHost 195.168.1.1 -searchTimeout 300')
```

- Using Jython list:

```
AdminTask.configureAppLDAPUserRegistry(['-securityDomainName', 'testDomain', '-ldapServerType', 'NETSCAPE', '-ldapHost', '195.168.1.1', '-searchTimeout', '300'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAppLDAPUserRegistry('-interactive')
```

configureAppLocalOSUserRegistry

The `configureAppLocalOSUserRegistry` command configures a local operating system user registry in a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Optional parameters

-realmName

Specifies the realm of the user registry. The system automatically generates a realm name if you do not specify a value for the `-realmName` parameter. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAppLocalOSUserRegistry('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.configureAppLocalOSUserRegistry(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAppLocalOSUserRegistry('-interactive')
```

configureAppWIMUserRegistry

The `configureAppWIMUserRegistry` command configures federated repository user registries in a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Optional parameters

-realmName

Specifies the realm of the user registry. The system automatically generates a realm name if you do not specify a value for the `-realmName` parameter. (String)

-verifyRegistry

Specifies whether to verify that the user registry configuration is correct. If you set this parameter to `true`, then the system verifies the registry by making a call to the user registry to verify the admin ID. If you specify a server ID and password, then the system verifies the user and password with the user registry. Set the parameter to `false` to store the attributes in the configuration without validation. The command verifies the registry configuration by default. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAppWIMUserRegistry('-securityDomainName testDomain -realmName testRealm')
```

- Using Jython list:

```
AdminTask.configureAppWIMUserRegistry(['securityDomainName', 'testDomain', '-realmName', 'testRealm'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAppWIMUserRegistry('-interactive')
```

getLTPATimeout

The `getLTPATimeout` command displays the number of seconds that the system waits before the LTPA request reaches timeout.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

Return value

The command returns the number of seconds that the server waits before the LTPA request is cancelled.

Batch mode example usage

- Using Jython string:

```
AdminTask.getLTPATimeout('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getLTPATimeout(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getLTPATimeout('-interactive')
```

setLTPATimeout

The setLTPATimeout command sets the amount of time that the system waits before the LTPA request becomes invalid.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

-timeout

Specifies the amount of time, in seconds, before the request times out. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.setLTPATimeout('-timeout 120')
```

- Using Jython list:

```
AdminTask.setLTPATimeout(['timeout', '120'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setLTPATimeout('-interactive')
```

getUserRegistryInfo

The `getUserRegistryInfo` command displays information about a user registry in a security domain or in the global security configuration. If you do not specify a value for the `-userRegistryType` parameter, the command returns the active user registry information.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

-userRegistryType

Specifies the type of user registry. Specify `LDAPUserRegistry` for LDAP user registries. Specify `WIMUserRegistry` for federated repository user registries. Specify `CustomUserRegistry` for custom user registries. Specify `LocalOSUserRegistry` for local operating system user registries. (String)

Return value

The command returns configuration information in the form of attribute and value pairs for the user registry object of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.getUserRegistryInfo('-securityDomainName testDomain -userRegistryType LDAPUserRegistry')
```

- Using Jython list:

```
AdminTask.getUserRegistryInfo(['securityDomainName', 'testDomain', '-userRegistryType', 'LDAPUserRegistry'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getUserRegistryInfo('-interactive')
```

unconfigureUserRegistry

The `unconfigureUserRegistry` command modifies the user registry. For a global security configuration, the command reduces the user registry to the minimum registry values. For application-level security, the command removes the user registry from the security domain of interest.

Target object

None.

Required parameters

-userRegistryType

Specifies the type of user registry. Specify `LDAPUserRegistry` for LDAP user registries. Specify `WIMUserRegistry` for federated repository user registries. Specify `CustomUserRegistry` for custom user registries. Specify `LocalOSUserRegistry` for local operating system user registries. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureUserRegistry('-userRegistryType WIMUserRegistry -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureUserRegistry(['userRegistryType', 'WIMUserRegistry', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureUserRegistry('-interactive')
```

configureLoginEntry

The `configureLoginEntry` command configures a Java Authentication and Authorization Service (JAAS) login entry in a security domain or in the global security configuration. You can use this command to modify existing JAAS login entries or to create new login entries.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify `system` for the system login type or `application` for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. If you do not specify a security domain name, the system updates the global security configuration. (String)

-loginModules

Specifies a comma (,) separated list of login module class names. Specify the list in the order that the system calls them. (String)

-authStrategies

Specifies a comma-separated list of authentication strategies that sets the authentication behavior as authentication proceeds down the list of login modules. You must specify one authentication strategy for each login module. (String)

Specify one or many of the following values in a comma (,) separated list:

- REQUIRED

Specifies that the LoginModule module is required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list for each realm.

- REQUISITE

Specifies that the LoginModule module is required to succeed. If authentication is successful, the process continues down the LoginModule list in the realm entry. If authentication fails, control immediately returns to the application. Authentication does not proceed down the LoginModule list.

- SUFFICIENT

Specifies that the LoginModule module is not required to succeed. If authentication succeeds, control immediately returns to the application. Authentication does not proceed down the LoginModule list. If authentication fails, the process continues down the list.

- OPTIONAL

Specifies that the LoginModule module is not required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list.

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureLoginEntry(['-loginType application -loginEntryAlias JAASLoginEntry1 -authStrategies "REQUIRED,REQUISITE"'])
```

- Using Jython list:

```
AdminTask.configureLoginEntry(['loginType', 'application', '-loginEntryAlias', 'JAASLoginEntry1', '-authStrategies', 'REQUIRED,REQUISITE'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureLoginEntry('-interactive')
```

configureLoginModule

The configureLoginModule command modifies an existing login module or creates a new login module on an existing JAAS login entry in the global security configuration or in a security domain.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify system for the system login type or application for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

-loginModule

Specifies the name of the login module. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. (String)

-useLoginModuleProxy

Specifies that the JAAS loads the login module proxy class. JAAS then delegates calls to the login module classes that are defined in the Module class name field. Specify true to use the login module proxy. (Boolean)

-authStrategy

Specifies the authentication behavior as authentication proceeds down the list of login modules.
(String)

Specify one of the following values:

- REQUIRED

Specifies that the LoginModule module is required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list for each realm.

- REQUISITE

Specifies that the LoginModule module is required to succeed. If authentication is successful, the process continues down the LoginModule list in the realm entry. If authentication fails, control immediately returns to the application. Authentication does not proceed down the LoginModule list.

- SUFFICIENT

Specifies that the LoginModule module is not required to succeed. If authentication succeeds, control immediately returns to the application. Authentication does not proceed down the LoginModule list. If authentication fails, the process continues down the list.

- OPTIONAL

Specifies that the LoginModule module is not required to succeed. Whether authentication succeeds or fails, the process still continues down the LoginModule list.

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: ["attr1=value1","attr2=value2"]
(String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureLoginModule('-loginType application -loginEntryAlias JAASLoginEntry1 -loginModule class1')
```

- Using Jython list:

```
AdminTask.configureLoginModule(['loginType', 'application', '-loginEntryAlias', 'JAASLoginEntry1', '-loginModule', 'class1'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureLoginModule('-interactive')
```

getJAASLoginEntryInfo

The getJAASLoginEntryInfo command displays configuration for a specific JAAS login entry.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify system for the system login type or application for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an attribute list that contains configuration information for the JAAS login entry of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.getJAASLoginEntryInfo('-loginType application -loginEntryAlias JAASLoginEntry -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getJAASLoginEntryInfo(['loginType', 'application', '-loginEntryAlias', 'JAASLoginEntry', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getJAASLoginEntryInfo('-interactive')
```

listJAASLoginEntries

The `listJAASLoginEntries` command displays each defined JAAS login modules for given type in a security domain or the global security configuration.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify `system` for the system login type or `application` for the application login type. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an array of attribute lists that contain the login entries for the login type of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.listJAASLoginEntries('-loginType application -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.listJAASLoginEntries(['loginType', 'application', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listJAASLoginEntries('-interactive')
```

listLoginModules

The listLoginModules command displays the class names and associated options for a specific JAAS login module in a security domain or in the global security configuration.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify *system* for the system login type or *application* for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

Return value

The command returns an array that contains the login modules in a specific login entry.

Batch mode example usage

- Using Jython string:

```
AdminTask.listLoginModules('-loginType system -loginEntryAlias JAASLoginEntry')
```

- Using Jython list:

```
AdminTask.listLoginModules(['loginType', 'system', '-loginEntryAlias', 'JAASLoginEntry'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listLoginModules('-interactive')
```

unconfigureJAASLoginEntry

The unconfigureJAASLoginEntry command removes a JAAS login entry from the global security configuration or a security domain. You cannot remove all login entries. The command returns an error if it cannot remove the login entry of interest.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify *system* for the system login type or *application* for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureJAASLoginEntry('-loginType application -loginEntryAlias myLoginEntry')
```

- Using Jython list:

```
AdminTask.unconfigureJAASLoginEntry(['loginType', 'application', '-loginEntryAlias', 'myLoginEntry'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureJAASLoginEntry('-interactive')
```

unconfigureLoginModule

The `unconfigureLoginModule` command removes a login module class from a login module entry.

Target object

None.

Required parameters

-loginType

Specifies the type of JAAS login entry of interest. Specify `system` for the system login type or `application` for the application login type. (String)

-loginEntryAlias

Specifies an alias that identifies the JAAS login entry in the configuration. (String)

-loginModule

Specifies the name of the login module class to remove from the configuration. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureLoginModule('-loginType system -loginEntryAlias systemLoginEntry -loginModule moduleClass')
```

- Using Jython list:

```
AdminTask.unconfigureLoginModule(['loginType', 'system', '-loginEntryAlias', 'systemLoginEntry', '-loginModule', 'moduleClass'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureLoginModule('-interactive')
```

createAuthDataEntry

The createAuthDataEntry command creates an authentication data entry for a J2EE Connector architecture (J2C) connector in the global security or security domain configuration.

Target object

None.

Required parameters

-alias

Specifies the name that uniquely identifies the authentication data entry. (String)

-user

Specifies the J2C authentication data user ID. (String)

-password

Specifies the password to use for the target enterprise information system (EIS). (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain configuration. The application server uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

-description

Specifies a description of the authentication data entry. (String)

Return value

The command returns the object name of the new authentication data entry object.

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuthDataEntry('-alias dataEntry1 -user userID -password userIDpw')
```

- Using Jython list:

```
AdminTask.createAuthDataEntry(['alias', 'dataEntry1', '-user', 'userID', '-password', 'userIDpw'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuthDataEntry('-interactive')
```

deleteAuthDataEntry

The deleteAuthDataEntry command removes an authentication data entry for a J2C connector in a global security or security domain configuration.

Target object

None.

Required parameters

-alias

Specifies the name that uniquely identifies the authentication data entry. (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain configuration. The application server uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuthDataEntry('-alias dataEntry1')
```

- Using Jython list:

```
AdminTask.deleteAuthDataEntry(['alias', 'dataEntry1'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuthDataEntry('-interactive')
```

getAuthDataEntry

The `getAuthDataEntry` command displays information about an authentication data entry for the J2C connector in the global security configuration or for a specific security domain.

Target object

None.

Required parameters

-alias

Specifies the name that uniquely identifies the authentication data entry. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an attribute list that contains the authentication data entry attributes and values.

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuthDataEntry('-alias authDataEntry1 -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getAuthDataEntry(['alias', 'authDataEntry1', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuthDataEntry('-interactive')
```

listAuthDataEntries

The `listAuthDataEntries` command displays each authentication data entry in the global security configuration or in a security domain.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an array of attribute lists for each authentication data entry.

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuthDataEntries('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.listAuthDataEntries(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuthDataEntries('-interactive')
```

modifyAuthDataEntry

The `modifyAuthDataEntry` command modifies an authentication data entry for a J2C connector in the global security or security domain configuration.

Target object

None.

Required parameters

-alias

Specifies the name that uniquely identifies the authentication data entry. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

-user

Specifies the J2C authentication data user ID. (String)

-password

Specifies the password to use for the target enterprise information system (EIS). (String)

-description

Specifies a description for the authentication data entry. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuthDataEntry('-alias dataEntry1 -user userID1 -password newPassword')
```

- Using Jython list:

```
AdminTask.modifyAuthDataEntry(['alias', 'dataEntry1', '-user', 'userID1', '-password', 'newPassword'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuthDataEntry('-interactive')
```

configureCSInbound

The `configureCSInbound` command configures CSIv2 inbound authentication on a security domain or on the global security configuration. When configuring CSI inbound authentication in a security domain for the first time that the CSI objects are copied from global security so that any changes to that configuration are applied.

Target object

None.

Optional parameters**-securityDomainName**

Specifies the name of the security configuration. If one is not provided the task will work on the global security user registry configuration. (String)

-messageLevelAuth

Specifies whether clients connecting to this server must specify a user ID and password. Specify `Never` to disable the user ID and password requirement. Specify `Supported` to accept a user ID and password. Specify `Required` to require a user ID and password. (String)

-supportedAuthMechList

Specifies the authentication mechanism to use. Specify `KRB5` for Kerberos authentication, `LTPA` for Lightweight Third-Party Authentication, `BasicAuth` for BasicAuth authentication, and `custom` to use your own authentication token implementation. You can specify more than one in a space-separated list. (String)

-clientCertAuth

Specifies whether a client that connects to the server must connect using an SSL certificate. Specify `Never` to allow clients to connect without SSL certificates. Specify `Supported` to accept clients connecting with and without SSL certificates. Specify `Required` to require clients to use SSL certificate. (String)

-transportLayer

Specifies the transport layer support level. Specify `Never` to disable transport layer support. Specify `Supported` to enable transport layer support. Specify `Required` to require transport layer support. (String)

-sslConfiguration

Specifies the SSL configuration alias to use for inbound transport. (String)

-enableIdentityAssertion

Specifies whether to enable identity assertion. When using the identity assertion authentication method, the security token generated is a <wsse:UsernameToken> element that contains a <wsse:Username> element. Specify true for the -enableIdentityAssertion parameter to enable identity assertion. (Boolean)

-trustedIdentities

Specifies a list of trusted server identities, separated by the pipe character (|). To specify a null value, set the value of the -trustedIdentities parameter as an empty string (""). (String)

-statefulSession

Specifies whether to enable a stateful session. Specify true to enable a stateful session. (Boolean)

-enableAttributePropagation

Specifies whether to enable security attribute propagation. Security attribute propagation allows the application server to transport authenticated Subject contents and security context information from one server to another in your configuration. Specify true to enable security attribute propagation. (Boolean)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureCSIIInbound(['-securityDomainName testDomain -messageLevelAuth Required  
-supportedAuthMechList "KRB5 LTPA"'])
```

- Using Jython list:

```
AdminTask.configureCSIIInbound(['-securityDomainName', 'testDomain', '-messageLevelAuth', 'Required',  
'-supportedAuthMechList', 'KRB5 LTPA'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureCSIIInbound('-interactive')
```

configureCSIOutbound

The configureCSIOutbound command configures the CSIv2 outbound authentication in a security domain or in the global security configuration. When configuring CSI Outbound in a security domain for the first time, the application server copies the CSI objects from global security. Then, the application server applies the changes to that configuration from the command.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. (String)

-enableAttributePropagation

Specifies whether to enable security attribute propagation. Security attribute propagation allows the application server to transport authenticated Subject contents and security context information from one server to another in your configuration. Specify true to enable security attribute propagation. (Boolean)

-enableIdentityAssertion

Specifies whether to enable identity assertion. When using the identity assertion authentication

method, the security token generated is a <wsse:UsernameToken> element that contains a <wsse:Username> element. Specify true for the -enableIdentityAssertion parameter to enable identity assertion. (Boolean)

-useServerIdentity

Specifies whether to use the server identity to establish trust with the target server. Specify true to use the server identity. (Boolean)

-trustedId

Specifies the trusted identity that the application server uses to establish trust with the target server. (String)

-trustedIdentityPassword

Specifies the password of the trusted server identity. (String)

-messageLevelAuth

Specifies whether clients connecting to this server must specify a user ID and password. Specify includeNever to disable the user ID and password requirement. Specify Supported to accept a user ID and password. Specify Required to require a user ID and password. (String)

-supportedAuthMechList

Specifies the authentication mechanism to use. Specify KRB5 for Kerberos authentication, LTPA for Lightweight Third-Party Authentication, BasicAuth for BasicAuth authentication, and custom to use your own authentication token implementation. You can specify more than one in a space-separated list. (String)

-clientCertAuth

Specifies whether a client that connects to the server must connect using an SSL certificate. Specify Never to allow clients to connect without SSL certificates. Specify Supported to accept clients connecting with and without SSL certificates. Specify Required to require clients to use SSL certificate. (String)

-transportLayer

Specifies the transport layer support level. Specify Never to disable transport layer support. Specify Supported to enable transport layer support. Specify Required to require transport layer support. (String)

-sslConfiguration

Specifies the SSL configuration alias to use for inbound transport. (String)

-statefulSession

Specifies whether to enable a stateful session. Specify true to enable a stateful session. (Boolean)

-enableOutboundMapping

Specifies whether to enable custom outbound identity mapping. Specify true to enable custom outbound identity mapping. (Boolean)

-trustedTargetRealms

Specifies a list of target realms to trust. Separate each realm name with the pipe character (|). (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureCSIOutbound('-securityDomainName testDomain -useServerIdentity true -messageAuthLevel Supported')
```

- Using Jython list:

```
AdminTask.configureCSIOutbound(['securityDomainName', 'testDomain', '-useServerIdentity', 'true', '-messageAuthLevel', 'Supported'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureCSIOutbound('-interactive')
```

getCSII inboundInfo

The `getCSII inboundInfo` command displays information about the Common Secure Interoperability (CSI) inbound settings for the global security configuration or for a security domain.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

-displayModel

Specifies the output format of the configuration information. Specify `true` to return an attribute list of the model. Specify `false` to display an attribute of the value used to create the object. (Boolean)

Return value

The command returns an attribute list of the attributes and values of the CSI inbound object.

Batch mode example usage

- Using Jython string:

```
AdminTask.getCSII inboundInfo('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getCSII inboundInfo(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getCSII inboundInfo('-interactive')
```

getCSIOutboundInfo

The `getCSIOutboundInfo` command displays information for the CSI outbound settings for the global security configuration or for a security domain.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

-displayModel

Specifies the output format of the configuration information. Specify `true` to return an attribute list of the model. Specify `false` to display an attribute of the value used to create the object. (Boolean)

Return value

The command returns an attribute list that contains the attributes and values of the CSI outbound configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.getCSIOutboundInfo('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getCSIOutboundInfo(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getCSIOutboundInfo('-interactive')
```

unconfigureCSIInbound

The unconfigureCSIInbound command removes the CSI inbound information from a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureCSIInbound('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureCSIInbound(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureCSIInbound('-interactive')
```

unconfigureCSIOutbound

The unconfigureCSIOutbound command removes the CSI outbound information from a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureCSIOutbound('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureCSIOutbound(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureCSIOutbound('-interactive')
```

configureInterceptor

The `configureInterceptor` command modifies an existing interceptor or creates an interceptor if one does not exist.

Target object

None.

Required parameters

-interceptor

Specifies the trust association interceptor class name. (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain. If you do not specify a security domain, the command assigns the global security configuration. (String)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: "attr1=value1","attr2=value2" (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureInterceptor('-interceptor com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus  
-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.configureInterceptor(['interceptor', 'com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus',  
-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureInterceptor('-interactive')
```


configureTrustAssociation

The `configureTrustAssociation` command enables or disables the trust association. If the security domain does not have a trust association defined, the application server copies each trust association and its interceptors from the global security configuration.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. (String)

-enable

Specifies whether to enable trust association to act as a reverse proxy server. (Boolean)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureTrustAssociation('-securityDomainName testDomain -enable true')
```

- Using Jython list:

```
AdminTask.configureTrustAssociation(['securityDomainName', 'testDomain', '-enable', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureTrustAssociation('-interactive')
```

getTrustAssociationInfo

The `getTrustAssociationInfo` command displays configuration information for trust association.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an attribute list that contains attributes and values for trust association.

Batch mode example usage

- Using Jython string:

```
AdminTask.getTrustAssociationInfo('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getTrustAssociationInfo(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getTrustAssociationInfo('-interactive')
```

listInterceptors

The listInterceptors command displays the trust association interceptors that are configured in the global security or security domain configuration.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

Return value

The command returns an array list of each interceptor and the associated custom properties.

Batch mode example usage

- Using Jython string:

```
AdminTask.listInterceptors('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.listInterceptors(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listInterceptors('-interactive')
```

unconfigureInterceptor

The unconfigureInterceptor command removes a trust association interceptor from the global security configuration or from a security domain.

Target object

None.

Required parameters

-interceptor

Specifies the trust association interceptor class name. (String)

Optional parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the -securityDomainName parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureInterceptor('-interceptor com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus  
-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureInterceptor(['interceptor', 'com.ibm.ws.security.web.TAMTrustAssociationInterceptorPlus',  
'-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureInterceptor('-interactive')
```

unconfigureTrustAssociation

The unconfigureTrustAssociation command removes the trust association object from a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureTrustAssociation('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureTrustAssociation(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureTrustAssociation('-interactive')
```

configureAuthzConfig

The configureAuthzConfig command configures an external Java Authorization Contract for Containers (JACC) authorization provider in a security domain or the global security configuration.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security configuration. (String)

-useJACCProvider

Specifies whether to use a JACC provider. Specify true to use a JACC provider. (Boolean)

-name

Specifies the name of the JACC provider to use. (String)

-description

Specifies a description of the JACC provider. (String)

-j2eePolicyImplClassName

Specifies the class name of an implementation class that represents the `javax.security.jacc.policy.provider` property according to the specification. (String)

-policyConfigurationFactoryImplClassName

Specifies the class name of an implementation class that represents the `javax.security.jacc.PolicyConfigurationFactory.provider` property. (String)

-roleConfigurationFactoryImplClassName

Specifies the class name of an implementation class that implements the `com.ibm.wsspi.security.authorization.RoleConfigurationFactory` interface. (String)

-requiresEJBArgumentsPolicyContextHandler

Specifies whether policy providers require the Enterprise JavaBeans arguments policy context handler to make access decisions. Specify `true` to enable this option. (Boolean)

-initializeJACCProviderClassName

Specifies the class name of an implementation class that implements the `com.ibm.wsspi.security.authorization.InitializeJACCProvider` interface. (String)

-supportsDynamicModuleUpdates

Specifies whether the provider supports dynamic changes to the Web modules. Specify `true` to enable this option. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureAuthzConfig(['-securityDomainName testDomain -useJACCProvider true -name testProvider -description "JACC provider for testing"'])
```

- Using Jython list:

```
AdminTask.configureAuthzConfig(['securityDomainName', 'testDomain', '-useJACCProvider', 'true', '-name', 'testProvider', '-description', 'JACC provider for testing'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureAuthzConfig('-interactive')
```

configureSingleSignon

The `configureSingleSignon` command configures a single sign-on object in global security.

Target object

None.

Optional parameters

-enable

Specifies whether to enable single sign-on. Specify `true` to enable single sign-on, or `false` to disable single sign-on. (Boolean)

-requiresSSL

Specifies whether single sign-on requests send through HTTPS. Specify `true` to enable this option. (Boolean)

-domainName

Specifies the domain name that contains a set of hosts to which the single sign-on applies. (String)

-interoperable

Specifies interoperability options. Specify `true` to send an interoperable cookie to the browser to support back-level servers. Specify `false` to disable the sending of interoperable cookies. (Boolean)

-attributePropagation

Specifies whether to enable inbound security attribute propagation. Specify `true` to enable Web inbound security attribution propagation. Specify `false` to use the single sign-on token to log in and recreate the Subject from the user registry. (Boolean)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureSingleSignon('-enable true -domainName mycompany.com')
```

- Using Jython list:

```
AdminTask.configureSingleSignon(['enable', 'true', '-domainName', 'mycompany.com'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.configureSingleSignon('-interactive')
```

getActiveSecuritySettings

The `getActiveSecuritySettings` command displays the active security settings for global security or a specific security domain.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security domain configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns the active security settings for the security domain of interest or the global security configuration, which includes the following settings:

- `cacheTimeout`
- `issuePermissionWarning`
- `activeAuthMechanism`
- `enforceJava2Security`

- appSecurityEnabled
- enableGlobalSecurity (global security only)
- adminPreferredAuthMech (global security only)
- activeAuthMechanism (global security only)
- activeUserRegistry
- enforceFineGrainedJCASecurity
- dynUpdateSSLConfig (global security only)
- useDomainQualifiedUserNames
- customProperties

Batch mode example usage

- Using Jython string:

```
AdminTask.getActiveSecuritySettings('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getActiveSecuritySettings(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getActiveSecuritySettings('-interactive')
```

getAuthzConfigInfo

The `getAuthzConfigInfo` command displays information about an external JACC authorization provider in a security domain or the global security configuration.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security domain configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command returns an attribute list that contains the attributes and values that are associated with the JACC authorization provider.

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuthzConfigInfo('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.getAuthzConfigInfo(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuthzConfigInfo('-interactive')
```

getSingleSignon

The `getSingleSignon` command displays configuration information about the single sign-on object as defined in the global security configuration.

Target object

None.

Optional parameters

None.

Return value

The command returns an attribute list that contains the attributes and values of the single sign-on configuration.

Batch mode example usage

- Using Jython:

```
AdminTask.getSingleSignon()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getSingleSignon('-interactive')
```

setAdminActiveSecuritySettings

The `setAdminActiveSecuritySettings` command sets the active security settings on the global security object.

Target object

None.

Optional parameters

-enableGlobalSecurity

Specifies whether to enable global security. Specify `true` to enable global security, or specify `false` to disable global security. (Boolean)

-cacheTimeout

Specifies the amount of time, in seconds, before authentication data becomes invalid. (Integer)

-issuePermissionWarning

Specifies whether to issue a warning during application installation if the application requires security permissions. Specify `true` to enable the warning notification, or specify `false` to disable the warning notification. (Boolean)

-enforceJava2Security

Specifies whether to enable Java Platform, Enterprise Edition (Java EE) security. Specify `true` to enable Java EE security permissions checking, or specify `false` to disable Java EE security. (Boolean)

-enforceFineGrainedJCA Security

Specifies whether to restrict application access. Specify `true` to restrict application access to sensitive Java EE Connector Architecture (JCA) mapping authentication data. (Boolean)

-appSecurityEnabled

Specifies whether to enable application-level security. Specify `true` to enable application level security, or specify `false` to disable application-level security. (Boolean)

-dynUpdateSSLConfig

Specifies whether to dynamically update SSL configuration changes. Specify `true` to update SSL configuration changes dynamically, or specify `false` to update the SSL configuration when the server starts. (Boolean)

-activeAuthMechanism

Specifies the active authentication mechanism. Specify `LTPA` for LTPA authentication, `KRB5` for Kerberos authentication, or `RSAToken` for RSA token authorization. (String)

-adminPreferredAuthMech

Specifies the preferred authentication mechanism. Specify `LTPA` for LTPA authentication, `KRB5` for Kerberos authentication, or `RSAToken` for RSA token authorization. (String)

-activeUserRegistry

Specifies the active user registry for the server. (String)

-useDomainQualifiedUserNames

Specifies the type of user name to use. Specify `true` to use domain qualified user names, or specify `false` to use the short name. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAdminActiveSecuritySettings('-enableGlobalSecurity true -cacheTimeout 300  
-enforceJava2Security true -appSecurityEnabled true')
```

- Using Jython list:

```
AdminTask.setAdminActiveSecuritySettings(['enableGlobalSecurity', 'true', '-cacheTimeout',  
'300', '-enforceJava2Security', 'true', '-appSecurityEnabled', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setAdminActiveSecuritySettings('-interactive')
```

setAppActiveSecuritySettings

The `setAppActiveSecuritySettings` command sets the active security settings on a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Optional parameters

-cacheTimeout

Specifies the amount of time, in seconds, before authentication data becomes invalid. (Integer)

-issuePermissionWarning

Specifies whether to issue a warning during application installation if the application requires security permissions. Specify `true` to enable the warning notification, or specify `false` to disable the warning notification. (Boolean)

-enforceJava2Security

Specifies whether to enable Java Platform, Enterprise Edition (Java EE) security. Specify `true` to enable Java EE security permissions checking, or specify `false` to disable Java EE security. (Boolean)

-enforceFineGrainedJCA Security

Specifies whether to restrict application access. Specify `true` to restrict application access to sensitive Java EE Connector Architecture (JCA) mapping authentication data. (Boolean)

-appSecurityEnabled

Specifies whether to enable application-level security. Specify `true` to enable application level security, or specify `false` to disable application-level security. (Boolean)

-activeUserRegistry

Specifies the active user registry for the server. (String)

-useDomainQualifiedUserNames

Specifies the type of user name to use. Specify `true` to use domain qualified user names, or specify `false` to use the short name. (Boolean)

-customProperties

Specifies a comma separated list of quoted attribute and value pairs that the system stores as custom properties on the user registry object. For example, use the format: `"attr1=value1","attr2=value2"` (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAppActiveSecuritySettings('-securityDomainName testDomain -issuePermissionWarning false  
-enforceFineGrainedJCA Security true')
```

- Using Jython list:

```
AdminTask.setAppActiveSecuritySettings(['securityDomainName', 'testDomain', '-issuePermissionWarning',  
'false', '-enforceFineGrainedJCA Security', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setAppActiveSecuritySettings('-interactive')
```

unconfigureAuthzConfig

The `unconfigureAuthzConfig` command removes an external JACC authorization provider from the global security configuration or a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureAuthzConfig('-securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureAuthzConfig(['securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unconfigureAuthzConfig('-interactive')
```

unsetAppActiveSecuritySettings

The `unsetAppActiveSecuritySettings` command removes an attribute from the global security configuration or a security domain.

Target object

None.

Required parameters

-securityDomainName

Specifies the name of the security configuration. The command uses the global security configuration if you do not specify a value for the `-securityDomainName` parameter. (String)

Optional parameters

-unsetAppSecurityEnabled

Specifies whether to remove the attribute that enables application security. Specify `true` to remove the attribute. (Boolean)

-unsetActiveUserRegistry

Specifies whether to remove the active user registry attribute. Specify `true` to remove the attribute. (Boolean)

-unsetUseDomainQualifiedUserNames

Specifies whether to remove the user domain qualified user names attribute. Specify `true` to remove the attribute. (Boolean)

-unsetEnforceJava2Security

Specifies whether to remove the Java EE security attribute. Specify `true` to remove the attribute. (Boolean)

-unsetEnforceFineGrainedJCASecurity

Specifies whether to remove the fine-grained JCA security attribute. Specify `true` to remove the attribute. (Boolean)

-unsetIssuePermissionWarning

Specifies whether to remove the attribute that issues user permission warnings. Specify `true` to remove the attribute. (Boolean)

-unsetCacheTimeout

Specifies whether to remove the cache timeout attribute. Specify `true` to remove the attribute. (Boolean)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unsetAppActiveSecuritySettings('-securityDomainName testDomain -unsetAppSecurityEnabled true -unsetPermissionWarning true')
```

- Using Jython list:

```
AdminTask.unsetAppActiveSecuritySettings(['securityDomainName', 'testDomain', '-unsetAppSecurityEnabled', 'true', '-unsetPermissionWarning', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.unsetAppActiveSecuritySettings('-interactive')
```

SecurityRealmInfoCommands command group for the AdminTask object

You can use the Jython scripting language to manage security realm configurations with the `wsadmin` tool. Use the commands and parameters in the `SecurityRealmInfoCommands` group to query and manage trusted realms.

Use the following commands to manage trusted realms in your security configuration:

- “`addTrustedRealms`”
- “`configureTrustedRealms`” on page 940
- “`listRegistryGroups`” on page 941
- “`listRegistryUsers`” on page 942
- “`listSecurityRealms`” on page 943
- “`listTrustedRealms`” on page 943
- “`removeTrustedRealms`” on page 944
- “`unconfigureTrustedRealms`” on page 945

addTrustedRealms

The `addTrustedRealms` command adds a realm or list of realms to the list of trusted realms for global security or in a security domain.

Target object

None.

Required parameters

-communicationType

Specifies whether to trusted realms to inbound or outbound communication. Specify `inbound` to configure inbound communication. Specify `outbound` to configure outbound communication. (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain of interest. If you do not specify a value for this parameter, the command uses the global security configuration. (String)

-realmList

Specifies a realm or list of realms to configure as trusted realms. (String)

Separate each realm in the list with the pipe character (|) as the following example demonstrates:

```
realm1|realm2|realm3
```

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.addTrustedRealms('-communicationType inbound -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.addTrustedRealms(['-communicationType', 'inbound', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.addTrustedRealms('-interactive')
```

configureTrustedRealms

The `configureTrustedRealms` command configures trusted realms. Use this command to replace the list of trusted realms and to clear each realm from the list. To add realms to the trusted realm list, use the `addInboundTrustedRealm` command.

Target object

None.

Required parameters

-communicationType

Specifies whether to configure the security domains, realms, or global security configuration for inbound or outbound communication. Specify `inbound` to configure inbound communication. Specify `outbound` to configure outbound communication. (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain of interest. If you do not specify a value for this parameter, the command uses the global security configuration. (String)

-realmList

Specifies a list of realms to configure as trusted realms. (String)

Separate each realm in the list with the pipe character (|) as the following example demonstrates:

```
realm1|realm2|realm3
```

-trustAllRealms

Specifies whether to trust all realms. Specify `true` to trust all realms. If you specify `true` for this parameter, the command does not use the `-realmList` parameter. (Boolean)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.configureTrustedRealms('-communicationType inbound -realmList realm1|realm2|realm3')
```

- Using Jython list:

```
AdminTask.configureTrustedRealms(['-communicationType', 'inbound', '-realmList', 'realm1|realm2|realm3'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.configureTrustedRealms('-interactive')
```

- Using Jython list:

listRegistryGroups

The `listRegistryGroups` command displays the groups in the user registry that belong to the security realm, security domain, or resource name of interest.

Target object

None.

Optional parameters

-securityRealmName

Specifies name of the security realm of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. (String)

-resourceName

Specifies the name of the resource of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. (String)

-securityDomainName

Specifies the name of the security domain of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. (String)

-displayAccessIds

Specifies whether to display the access IDs for each group. Specify `true` to display the access ID and group name for each group that the command returns. (Boolean)

-groupFilter

Specifies a filter that the command uses to query for groups. For example, specify `test*` to return groups that begin with the `test` string. By default, the command returns all groups. (String)

-numberOfGroups

Specifies the number of groups to return. The default number of groups that the command displays is 20. (Integer)

Return value

The command returns an array of group names. If you specified the `-displayAccessId` parameter, the command returns an array of attribute lists which contain the group name and group access ID.

Batch mode example usage

- Using Jython string:

```
AdminTask.listRegistryGroups('-securityDomainName myTestDomain -groupFilter test* -numberOfGroups 10')
```

- Using Jython list:

```
AdminTask.listRegistryGroups(['-securityDomainName', 'myTestDomain', '-groupFilter', 'test*', '-numberOfGroups', '10'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listRegistryGroups('-interactive')
```

listRegistryUsers

The `listRegistryUsers` command displays the users in the user registry for a specific security realm, resource name, or domain name.

Target object

None.

Optional parameters

-securityDomainName

Specifies the name of the security domain of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. If you do not specify the `securityDomainName`, `resourceName`, or `securityRealmName` parameter, the system uses the active user registry from the global security configuration. (String)

-resourceName

Specifies the name of the resource of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. If you do not specify the `securityDomainName`, `resourceName`, or `securityRealmName` parameter, the system uses the active user registry from the global security configuration. (String)

-securityRealmName

Specifies the name of the security realm of interest. The `securityDomainName`, `resourceName`, and `securityRealmName` parameters are mutually exclusive. Do not specify more than one of these parameters. If you do not specify the `securityDomainName`, `resourceName`, or `securityRealmName` parameter, the system uses the active user registry from the global security configuration. (String)

-displayAccessIds

Specifies whether to display the access IDs for each group. Specify `true` to display the access ID and group name for each group that the command returns. (Boolean)

-userFilter

Specifies the filter that the command uses to query for users. For example, specify `test*` to display each user name that starts with the `test` string. By default, the command returns all users. (String)

-numberOfUsers

Specifies the number of users to return. The default number of groups that the command displays is 20. (Integer)

Return value

The command returns an array of user names. If you specify the `-displayAccessId` parameter, the command returns an array of attribute lists that contain the user ID and user access IDs.

Batch mode example usage

- Using Jython string:

```
AdminTask.listRegistryUsers('-securityRealmName defaultWIMFileBasedRealm -displayAccessIds true')
```

- Using Jython list:


```
AdminTask.listRegistryUsers(['-securityRealmName', 'defaultWIMFileBasedRealm', '-displayAccessIds', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listRegistryUsers('-interactive')
```

listSecurityRealms

The `listSecurityRealms` command displays each security realm from global security configuration and the security domains.

Target object

None.

Return value

The command returns an array of realm names.

Batch mode example usage

- Using Jython string:

```
AdminTask.listSecurityRealms()
```

- Using Jython list:

```
AdminTask.listSecurityRealms()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listSecurityRealms('-interactive')
```

listTrustedRealms

The `listTrustedRealms` command displays a list of trusted realms for a security domain, resource, or realm. If you do not specify a security domain, resource name, or realm name, then the command returns a list of trusted realms from the global security configuration. The `securityRealmName`, `resourceName`, and `securityDomainName` parameters are mutually exclusive.

Target object

None.

Required parameters

-communicationType

Specifies whether to list the trusted realms for inbound or outbound communication. Specify `inbound` to configure inbound communication. Specify `outbound` to configure outbound communication. (String)

Optional parameters

-securityRealmName

Specifies name of the security realm of interest. If you use this parameter, do not use the `resourceName` or `securityDomainName` parameters. (String)

-resourceName

Specifies the name of the resource of interest. If you use this parameter, do not use the `securityRealmName` or `securityDomainName` parameters. (String)

-securityDomainName

Specifies the name of the security domain of interest. If you use this parameter, do not use the `resourceName` or `securityRealmName` parameters. (String)

-expandRealmList

Specifies whether to return each realm name when the `trustAllRealms` property is enabled. Specify `true` to return each realm name. Specify `false` to return the `trustAllRealms` property. (Boolean)

-includeCurrentRealm

Specifies whether to include the current realm in the list of trusted realms. Specify `true` to include the current realm, or specify `false` to exclude the current realm from the list of trusted realms. (Boolean)

Return value

The command returns an array of trusted realm names. If the realm, resource, or security domain of interest is configured to trust all realms, the command returns the `trustAllRealms` string.

Batch mode example usage

- Using Jython string:

```
AdminTask.listTrustedRealms('-communicationType inbound -resourceName myApplication')
```

- Using Jython list:

```
AdminTask.listTrustedRealms(['-communicationType', 'inbound', '-resourceName', 'myApplication'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listTrustedRealms('-interactive')
```

removeTrustedRealms

The `removeTrustedRealms` command removes realms from a trusted realm list in a security domain or in the global security configuration.

Target object

None.

Required parameters

-communicationType

Specifies whether to remove trusted realms from inbound or outbound communication. Specify `inbound` to configure inbound communication. Specify `outbound` to configure outbound communication. (String)

-realmList

Specifies a list of realms to remove from trusted realms. (String)

Separate each realm in the list with the pipe character (`|`) as the following example demonstrates:

```
realm1|realm2|realm3
```

Optional parameters

-securityDomainName

Specifies the name of the security domain of interest. If you do not specify a security domain, the command uses the global security configuration. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.removeTrustedRealms('-communicationType inbound -realmList realm1|realm2|realm3')
```

- Using Jython list:

```
AdminTask.removeTrustedRealms(['-communicationType inbound -realmList realm1|realm2|realm3'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.removeTrustedRealms('-interactive')
```

unconfigureTrustedRealms

The `unconfigureTrustedRealms` command removes the trusted realm object from the configuration.

Target object

None.

Required parameters

-communicationType

Specifies whether to unconfigure the trusted realms for inbound or outbound communication. Specify `inbound` to remove inbound communication configurations. Specify `outbound` to remove outbound communication configurations. (String)

Optional parameters

-securityDomainName

Specifies the name of the security domain of interest. If you do not specify a security domain, the command uses the global security configuration. (String)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.unconfigureTrustedRealms('-communicationType inbound -securityDomainName testDomain')
```

- Using Jython list:

```
AdminTask.unconfigureTrustedRealms(['-communicationType', 'inbound', '-securityDomainName', 'testDomain'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.unconfigureTrustedRealms('-interactive')
```

NamingAuthzCommands command group for the AdminTask object

You can use the Jython scripting language to configure naming roles for groups and users with the `wsadmin` tool. Use the commands and parameters in the `NamingAuthzCommands` group to assign, remove, and query naming role configuration. `CosNaming` security offers increased granularity of security control over `CosNaming` functions.

A number of naming roles are defined to provide the degrees of authority that are needed to perform certain application server naming service functions. The authorization policy is only enforced when global security is enabled.

Use the following commands to manage the naming service functions:

- `listGroupsForNamingRoles`
- `listUsersForNamingRoles`
- `mapGroupsToNamingRole`
- `mapUsersToNamingRole`
- `removeGroupsFromNamingRole`
- `removeUsersFromNamingRole`

listGroupsForNamingRoles

The `listGroupsForNamingRoles` command displays the groups and special subjects that are mapped to the naming roles.

Target object

None.

Return value

The command returns a list of the groups and special subjects associated with each naming role.

Batch mode example usage

- Using Jython:

```
AdminTask.listGroupsForNamingRoles()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listGroupsForNamingRoles('-interactive')
```

listUsersForNamingRoles

The `listUsersForNamingRoles` command displays the users that are mapped to the naming roles.

Target object

None.

Return value

The command returns a list of the users associated with each naming role.

Batch mode example usage

- Using Jython:

```
AdminTask.listUsersForNamingRoles()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listUsersForNamingRoles('-interactive')
```

mapGroupsToNamingRole

The `mapGroupsToNamingRole` command maps groups, special subjects, or groups and special subjects to the naming roles.

Target object

None.

Required parameters

-roleName

Specifies the name of the naming role. (String)

Four name space security roles are available: `CosNamingRead`, `CosNamingWrite`, `CosNamingCreate`, and `CosNamingDelete`. The roles have authority levels from low to high, as the following table defines:

Role name	Description
<code>CosNamingRead</code>	You can query the application server name space using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The <code>EVERYONE</code> special-subject is the default policy for this role.
<code>CosNamingWrite</code>	You can perform write operations such as JNDI bind, rebind, or unbind, and <code>CosNamingRead</code> operations.
<code>CosNamingCreate</code>	You can create new objects in the name space through operations such as JNDI <code>createSubcontext</code> and <code>CosNamingWrite</code> operations.
<code>CosNamingDelete</code>	You can destroy objects in the name space, for example using the JNDI <code>destroySubcontext</code> method and <code>CosNamingCreate</code> operations.

Optional parameters

-groupids

Specifies the names of the groups to map to the naming roles. (String[])

-accessids

Specifies the access IDs of the users in the format `<group:realmName/uniqueID>`. (String[])

-specialSubjects

Specifies the special subjects to map. (String[])

The special subjects include `EVERYONE`, `ALLAUTHENTICATED`, `ALLAUTHENTICATEDINTRUSTEDREALMS`, as the following table defines:

Header	Header
<code>EVERYONE</code>	Maps 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.
<code>ALLAUTHENTICATED</code>	Maps each authenticated user to a specified role. When you map each authenticated user to a specified role, each valid user in the current registry who has been authenticated can access resources that are protected by this role.
<code>ALLAUTHENTICATEDINTRUSTEDREALMS</code>	Maps each authenticated user to a specified role. When you map each authenticated user to a specified role, each valid user in the current registry who has been authenticated can access resources that are protected by this role in the trusted realm.

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.mapGroupsToNamingRole(['-roleName CosNamingCreate -groupids [group1, group2]'])
```

- Using Jython list:

```
AdminTask.mapGroupsToNamingRole(['-roleName', 'CosNamingCreate', '-groupids', '[group1, group2]'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.mapGroupsToNamingRole('-interactive')
```

mapUsersToNamingRole

The `mapUsersToNamingRole` command maps users to the naming roles.

Target object

None.

Required parameters

-roleName

Specifies the name of the naming role. (String)

Four name space security roles are available: `CosNamingRead`, `CosNamingWrite`, `CosNamingCreate`, and `CosNamingDelete`. The roles have authority levels from low to high, as the following table defines:

Role name	Description
<code>CosNamingRead</code>	You can query the application server name space using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The <code>EVERYONE</code> special-subject is the default policy for this role.
<code>CosNamingWrite</code>	You can perform write operations such as JNDI bind, rebind, or unbind, and <code>CosNamingRead</code> operations.
<code>CosNamingCreate</code>	You can create new objects in the name space through operations such as JNDI <code>createSubcontext</code> and <code>CosNamingWrite</code> operations.
<code>CosNamingDelete</code>	You can destroy objects in the name space, for example using the JNDI <code>destroySubcontext</code> method and <code>CosNamingCreate</code> operations.

Optional parameters

-userids

Specifies the user IDs to map to the naming roles of interest. (String[])

-accessids

Specifies the access IDs of the users in the format `<user:realmName/uniqueID>`. (String[])

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.mapUsersToNamingRole([-roleName CosNamingDelete -userids [user1, user2, user3]])
```

- Using Jython list:

```
AdminTask.mapUsersToNamingRole(['-roleName', 'CosNamingDelete', '-userids', '[user1, user2, user3]'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.mapUsersToNamingRole('-interactive')
```

removeGroupsFromNamingRole

The `removeGroupsFromNamingRole` command removes groups, special subjects, or groups and special subjects from a naming role.

Target object

None.

Required parameters

-roleName

Specifies the name of the naming role. (String)

Four name space security roles are available: `CosNamingRead`, `CosNamingWrite`, `CosNamingCreate`, and `CosNamingDelete`. The roles have authority levels from low to high, as the following table defines:

Role name	Description
<code>CosNamingRead</code>	You can query the application server name space using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The <code>EVERYONE</code> special-subject is the default policy for this role.
<code>CosNamingWrite</code>	You can perform write operations such as JNDI bind, rebind, or unbind, and <code>CosNamingRead</code> operations.
<code>CosNamingCreate</code>	You can create new objects in the name space through operations such as JNDI <code>createSubcontext</code> and <code>CosNamingWrite</code> operations.
<code>CosNamingDelete</code>	You can destroy objects in the name space, for example using the JNDI <code>destroySubcontext</code> method and <code>CosNamingCreate</code> operations.

Optional parameters

-groupids

Specifies the names of the groups to remove from the naming roles of interest. (String[])

-specialSubjects

Specifies the special subjects to remove. (String[])

The special subjects include `EVERYONE`, `ALLAUTHENTICATED`, `ALLAUTHENTICATEDINTRUSTEDREALMS`, as the following table defines:

Header	Header
<code>EVERYONE</code>	Maps 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.
<code>ALLAUTHENTICATED</code>	Maps each authenticated user to a specified role. When you map each authenticated user to a specified role, each valid user in the current registry who has been authenticated can access resources that are protected by this role.
<code>ALLAUTHENTICATEDINTRUSTEDREALMS</code>	Maps each authenticated user to a specified role. When you map each authenticated user to a specified role, each valid user in the current registry who has been authenticated can access resources that are protected by this role in the trusted realm.

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.removeGroupsFromNamingRole('-roleName CosNamingRead -groupids [group1, group2] -specialSubjects EVERYONE')
```

- Using Jython list:

```
AdminTask.removeGroupsFromNamingRole(['-roleName', 'CosNamingRead', '-groupids', '[group1, group2]', '-specialSubjects', 'EVERYONE'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.removeGroupsFromNamingRole('-interactive')
```

removeUsersFromNamingRole

The `removeUsersFromNamingRole` command removes users from a naming role.

Target object

None.

Required parameters

-roleName

Specifies the name of the naming role. (String)

Four name space security roles are available: `CosNamingRead`, `CosNamingWrite`, `CosNamingCreate`, and `CosNamingDelete`. The roles have authority levels from low to high, as the following table defines:

Role name	Description
<code>CosNamingRead</code>	You can query the application server name space using, for example, the Java Naming and Directory Interface (JNDI) lookup method. The <code>EVERYONE</code> special-subject is the default policy for this role.
<code>CosNamingWrite</code>	You can perform write operations such as JNDI bind, rebind, or unbind, and <code>CosNamingRead</code> operations.
<code>CosNamingCreate</code>	You can create new objects in the name space through operations such as JNDI <code>createSubcontext</code> and <code>CosNamingWrite</code> operations.
<code>CosNamingDelete</code>	You can destroy objects in the name space, for example using the JNDI <code>destroySubcontext</code> method and <code>CosNamingCreate</code> operations.

Optional parameters

-userids

Specifies the user IDs to remove from the naming roles of interest. (String[])

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.removeUsersFromNamingRole('-roleName CosNamingRead')
```

- Using Jython list:

```
AdminTask.removeUsersFromNamingRole(['-roleName', 'CosNamingRead'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.removeUsersFromNamingRole('-interactive')
```

Utility scripts

The scripting library provides multiple script procedures to automate your application configurations. This topic provides usage information for scripts that set notification options, save configuration changes, and display scripting library information.

Each utility script procedure is located in the `app_server_root/scriptLibraries/utilities/V70` directory. Use the following script procedures to perform utility functions:

- “convertToList” on page 951
- “debugNotice” on page 951
- “getExceptionText” on page 951
- “fail” on page 952
- “fileSearch” on page 952
- “getResourceBundle” on page 953
- “getScriptLibraryFiles” on page 953
- “getScriptLibraryList” on page 953

- “getScriptLibraryPath” on page 953
- “help” on page 953
- “infoNotice” on page 954
- “save” on page 954
- “setDebugNotices” on page 954
- “setFailOnErrorDefault” on page 954
- “sleepDelay” on page 955
- “warningNotice” on page 955

convertToList

This script converts a string to a list. For example, the `AdminApp.list()` command returns a string of application names. Use the `convertToList` script to change the output to a list format, such as `['DefaultApplication', 'a1', 'a2', 'ivtApp', 'query']`.

To run the script, run a command that returns a string output and set the output to a variable, as defined in the following table:

Argument	Description
<i>variable</i>	Specifies the name of the variable that contains the string to convert to a list.

Syntax

```
AdminUtilities.convertToList(variable)
```

Example usage

```
apps=AdminApp.list()
AdminUtilities.convertToList(apps)
```

debugNotice

This script sets the debug notice text.

To run the script, specify the message argument, as defined in the following table:

Argument	Description
<i>message</i>	Specifies the message text for the debug notice.

Syntax

```
AdminUtilities.debugNotice(message)
```

Example usage

```
AdminUtilities.debugNotice("Server is started")
```

getExceptionText

This script displays the exception message for a specific exception type, exception value, or traceback information.

To run the script, specify the type, value, or traceback arguments, as defined in the following table:

Argument	Description
<i>type</i>	Specifies the exception type of interest. The exception type represents the class object of the exception.

Argument	Description
<i>value</i>	Specifies the exception value of interest. The value represents the instance object that is the argument of the exception or the second argument of the raise statement.
<i>traceback</i>	Specifies the traceback information of interest. The traceback object contains special attributes, including the line number where the error occurred. Do not assign <i>traceback</i> to a local variable in the function that handles the exception, as this assignment creates a circular reference.

Syntax

```
AdminUtilities.getExceptionText(type, value, traceback)
```

Example usage

```
AdminUtilities.getExceptionText("com.ibm.ws.scripting.ScriptingException",
"com.ibm.ws.scripting.ScriptingException: AdminControl service not available",
"")
```

fail

This script sets the failure message.

To run the script, specify the message argument, as defined in the following table:

Argument	Description
<i>message</i>	Specifies the message text for the failure notice.

Syntax

```
AdminUtilities.fail(message)
```

Example usage

```
AdminUtilities.fail("The script failed")
```

fileSearch

This script searches the file system based on a specific path or directory.

To run the script, specify the path or directory arguments, as defined in the following table:

Argument	Description
<i>path</i>	Specifies the file path to search for a specific file.
<i>directory</i>	Specifies the directory to search for a specific file.

Syntax

```
AdminUtilities.fileSearch(path, directory)
```

Example usage Windows

```
Paths = []
Directory = java.io.File("c:\\WebSphere\\AppServer\\scriptLibraries")
AdminUtilities.fileSearch(directory, paths)
```

Linux
AIX
HP-UX
Solaris

```
Paths = []
Directory = java.io.File("//WebSphere/AppServer/scriptLibraries")
AdminUtilities.fileSearch(directory, paths)
```

getResourceBundle

This script displays an instance for the resource bundle of interest.

To run the script, specify the bundle name argument, as defined in the following table:

Argument	Description
<i>bundleName</i>	Specifies the name of the bundle of interest. For example, to get a message object from the ScriptingLibraryMessage resource bundle, specify <code>com.ibm.ws.scripting.resources.scriptLibraryMessage</code> .

Syntax

```
AdminUtilities.getResourceBundle(bundleName)
```

Example usage

```
AdminUtilities.getResourceBundle("com.ibm.ws.scripting.resources.scriptLibraryMessage")
```

getScriptLibraryFiles

This script displays the file path and file names for each script library file.

Syntax

```
AdminUtilities.getScriptLibraryFiles()
```

Example usage

```
AdminUtilities.getScriptLibraryFiles()
```

getScriptLibraryList

This script displays each script name in the script library.

Syntax

```
AdminUtilities.getScriptLibraryList()
```

Example usage

```
AdminUtilities.getScriptLibraryList()
```

getScriptLibraryPath

This script displays the file path to get to the script library files on your file system.

Syntax

```
AdminUtilities.getScriptLibraryPath()
```

Example usage

```
AdminUtilities.getScriptLibraryPath()
```

help

This script displays help information for the AdminUtilities script library, including general library information, script names, and script descriptions.

To run the script, optionally specify the name of the script of interest, as defined in the following table:

Argument	Description
<i>scriptName</i>	Optionally specifies the name of the AdminUtilities script of interest.

Syntax

```
AdminUtilities.help(scriptName)
```

Example usage

```
AdminUtilities.help("sleepDelay")
```

infoNotice

This script sets the text for the information notice of a command or script.

To run the script, specify the message argument, as defined in the following table:

Argument	Description
<i>message</i>	Specifies the message text or a message ID such as "Application is installed" or <code>resourceBundle.getString("WASX7115I")</code> .

Syntax

```
AdminUtilities.infoNotice(message)
```

Example usage

```
AdminUtilities.infoNotice(resourceBundle.getString("WASX7115I"))
```

save

This script saves the configuration changes to your system.

Syntax

```
AdminUtilities.save()
```

Example usage

```
AdminUtilities.save()
```

setDebugNotices

This script enables and disables debug notices.

To run the script, specify the debug argument, as defined in the following table:

Argument	Description
<i>debug</i>	Specifies whether to enable or disable debug notices. Specify <code>true</code> to enable debug notices, or <code>false</code> to disable debug notices.

Syntax

```
AdminUtilities.setDebugNotices(debug)
```

Example usage

```
AdminUtilities.setDebugNotices("true")
```

setFailOnErrorDefault

This script enables or disables the fail on error behavior.

To run the script, specify the fail on error argument, as defined in the following table:

Argument	Description
<i>failOnError</i>	Specifies whether to enable or disable the fail on error behavior. Specify <code>true</code> to enable the fail on error behavior, or <code>false</code> to disable the behavior.

Syntax

```
AdminUtilities.setFailOnErrorDefault(failOnError)
```

Example usage

```
AdminUtilities.setFailOnErrorDefault("false")
```

sleepDelay

This script sets the number of seconds that the system waits for completion during two operations.

To run the script, specify the delay seconds argument, as defined in the following table:

Argument	Description
<i>delaySeconds</i>	Specifies the number of seconds to wait for completion.

Syntax

```
AdminUtilities.sleepDelay(delaySeconds)
```

Example usage

```
AdminUtilities.sleepDelay("10")
```

warningNotice

This script sets the text to display as the warning message.

To run the script, specify the message argument, as defined in the following table:

Argument	Description
<i>message</i>	Specifies the non-translated text for the warning notice or a message ID such as <code>resourceBundle.getString("WASX7411W")</code> .

Syntax

```
AdminUtilities.warningNotice(message)
```

Example usage

```
AdminUtilities.warningNotice(resourceBundle.getString("WASX7411W"))
```

Configuring the JACC provider for Tivoli Access Manager using the wsadmin utility

You can use the `wsadmin` utility to configure Tivoli Access Manager security for WebSphere Application Server.

About this task

Verify that all the managed servers, including node agents, are started. The following configuration is performed once on the deployment manager server. The configuration parameters are forwarded to managed servers, including node agents, when a synchronization is performed. The managed servers require their own restart for the configuration changes to take effect.

1. Start WebSphere Application Server.

2. Start the **wsadmin** command-line utility.
Run the **wsadmin** command from the *app_server_root/bin* directory.
3. At the **wsadmin** prompt, enter the following command:

```
$AdminTask configureTAM -interactive
```

You are prompted to enter the following information:

Option	Description
WebSphere Application Server node name	Specify a single node or enter an asterisk (*) to choose all nodes including the deployment manager.
Tivoli Access Manager Policy Server	Enter the name of the Tivoli Access Manager policy server and the connection port. Use the format, <i>policy_server : port</i> . The policy server communication port is set at the time of Tivoli Access Manager configuration. The default port is 7135.
Tivoli Access Manager Authorization Server	Enter the name of the Tivoli Access Manager authorization server. Use the format <i>auth_server : port : priority</i> . The authorization server communication port is set at the time of Tivoli Access Manager configuration. The default port is 7136. More than one authorization server can be specified by separating the entries with commas. Having more than one authorization server configured is useful for failover and performance. The priority value is the order of authorization server use. For example: <i>auth_server1:7136:1,auth_server2:7137:2</i> . A priority of 1 is still required when configuring against a single authorization server.
WebSphere Application Server administrator's distinguished name	Enter the full distinguished name of the WebSphere Application Server security administrator ID, as created in the "Creating the security administrative user" topic in the <i>Securing applications and their environment</i> PDF. For example: <i>cn=wasadmin,o=organization,c=country</i>
Tivoli Access Manager user registry distinguished name suffix	For example: <i>o=organization,c=country</i>
Tivoli Access Manager administrator's user name	Enter the Tivoli Access Manager administration user ID, as created at the time of Tivoli Access Manager configuration. This ID is usually, <i>sec_master</i> .
Tivoli Access Manager administrator's user password	Enter the password for the Tivoli Access Manager administrator.
Tivoli Access Manager security domain	Enter the name of the Tivoli Access Manager security domain that is used to store users and groups. If a security domain is not already established at the time of Tivoli Access Manager configuration, click Return to accept the default.
Embedded Tivoli Access Manager listening port set	WebSphere Application Server needs to listen on a TCP/IP port for authorization database updates from the policy server. More than one process can run on a particular node and machine so a list of ports is required for the processes. Enter the ports that are used as listening ports by Tivoli Access Manager clients, separated by a comma. If you specify a range of ports, separate the lower and higher values by a colon. For example, 7999, 9990:9999.

Option	Description
Defer	Set to <i>yes</i> , this option defers the configuration of the management server until the next restart. Set to <i>no</i> , configuration of the management server occurs immediately. Managed servers are configured on their next restart.

- When all information is entered, select **F** to save the configuration properties or **C** to cancel from the configuration process and discard entered information.

What to do next

Now enable the JACC provider for Tivoli Access Manager - see Enabling the JACC provider for Tivoli Access Manager topic in the *Securing applications and their environment* PDF.

Securing communications using the wsadmin tool

The application server provides several methods to secure communication between a server and a client. Use this topic to configure Secure Sockets Layer (SSL), keystores, certificate authorities, key sets and groups, and certificates.

- Configure secure communications using SSL.
 - Use the `SSLConfigCommands`, `SSLConfigGroupCommands`, `DynamicSSLConfigSelections` and `SSLTransport` command groups for the AdminTask object, and complete the following tasks to create and administer SSL configurations:
 - Create an SSL configuration at the node scope using scripting.
 - Automate SSL configurations using scripting.
- Create a keystore configuration.
 - Use the `KeyStoreCommands` command group for the AdminTask object, and complete the following tasks to create and administer keystore configurations.
 - Update default key store passwords using scripting
- Create a certificate authority (CA) client configuration.
 - A CA client object contains all of the configuration information necessary to connect to a third-party CA server. Use the `CAClientCommands` command group for the AdminTask object, and complete the following tasks to create and administer CA client objects in your configuration:
 - Configure CA clients using scripting
 - Administer CA clients using scripting
- Administer certificate configurations.
 - Use the `CertificateRequestCommands`, `PersonalCertificateCommands`, and `SignerCertificateCommands` command groups for the AdminTask object, and complete the following tasks to administer personal certificates, CA certificates, and self-signed certificates:
 - Create self-signed certificates using scripting
 - Configure CA certificates using scripting
- Create key sets and key groups.
 - Use the `KeySetCommands`, `KeySetGroupCommands`, and `KeyReferenceCommands` command groups for the AdminTask object to create and administer key set and group configurations.

Creating an SSL configuration at the node scope using scripting

An Secure Socket Layer (SSL) configuration references many other configuration objects. To help you make valid selections for the new SSL configuration before you create it, view information about existing configuration objects. Information about existing objects is also useful when you create a node scoped SSL configuration using the **createSSLConfig** command of the AdminTask object.

Before you begin

Before starting this task, the wsadmin tool must be running. See the Starting the wsadmin scripting client article for more information.

About this task

To use the information in this task effectively, familiarize yourself with the instructions in the Creating a Secure Sockets Layer configuration topic in the *Securing applications and their environment* PDF. Perform the following task to create an Secure Socket Layer (SSL) configuration at the node scope:

1. List the existing configuration objects. Perform any of the following:
 - List some of the configuration objects that you may need when you create a new SSL configuration. For example, you want to see which management scopes have already been defined. If the one you need does not exist you will need to create it.

– Using Jacl:

```
$AdminTask listManagementScopes {-scopeName (cell):BIRKT40Cell102:(node):BIRKT40Node02}
```

– Using Jython:

```
AdminTask.listManagementScopes ('[-scopeName (cell):BIRKT40Cell102:(node):BIRKT40Node02]')
```

This shows an existing cell scope and existing node scope that you can use. If you want to create a different scope, use the **createManagementScope** command of the AdminTask object to define a different one. The valid scope parameters are cell, nodegroup, node, server, cluster, and endpoint. See the Central management of Secure Sockets Layer configurations topic in the *Securing applications and their environment* PDF for more information on scope definitions.

- List the key stores that exist in the configuration including key stores and trust stores.

– Using Jacl:

```
$AdminTask listKeyStores -all true
```

– Using Jython:

```
AdminTask.listKeyStores('-all true')
```

Example output:

```
CellDefaultKeyStore(cells/BIRKT40Cell102|security.xml#KeyStore_1)
CellDefaultTrustStore(cells/BIRKT40Cell102|security.xml#KeyStore_2)
CellLTPAKeys(cells/BIRKT40Cell102|security.xml#KeyStore_3)
```

The previous example only lists the key stores for the default management scope which is also known as the cell scope. To obtain key stores for other scopes, specify the scopeName parameter, for example:

– Using Jacl:

```
$AdminTask listKeyStores {-scopeName (cell):BIRKT40Cell102:(node):BIRKT40Node02 }
```

– Using Jython:

```
$AdminTask listKeyStores ('[-scopeName (cell):BIRKT40Cell102:(node):BIRKT40Node02]')
```

Example output:

```
CellDefaultKeyStore(cells/BIRKT40Cell102|security.xml#KeyStore_1)
CellDefaultTrustStore(cells/BIRKT40Cell102|security.xml#KeyStore_2)
CellLTPAKeys(cells/BIRKT40Cell102|security.xml#KeyStore_3)
NodeDefaultKeyStore(cells/BIRKT40Cell102|security.xml#KeyStore_1134610924357)
NodeDefaultTrustStore(cells/BIRKT40Cell102|security.xml#KeyStore_1134610924377)
```

- List specific trust or key managers. Be sure to display the object name for the trust managers. You will need the object name for the SSL configuration because you can specify multiple trust manager instances.

– Using Jacl:

```
$AdminTask listTrustManagers {-scopeName (cell):BIRKT40Cell102:(node):BIRKT40Node02 -displayObjectName true }
```

- Using Jython:

```
AdminTask.listTrustManagers ('[-scopeName (cell):BIRKT40Cell02:(node):BIRKT40Node02 -displayObjectName true]')
```

Example output:

```
IbmX509(cells/BIRKT40Cell02|security.xml#TrustManager_1)
IbmPKIX(cells/BIRKT40Cell02|security.xml#TrustManager_2)
IbmX509(cells/BIRKT40Cell02|security.xml#TrustManager_1134610924357)
IbmPKIX(cells/BIRKT40Cell02|security.xml#TrustManager_1134610924377)
```

2. Create the node-scoped SSL configuration in interactive mode. Now that we have the information we need to choose from, we need to decide if these objects are sufficient or if we need to create new ones. For now, we will reuse what we've already got in the configuration and save creating new instances to task documents specific to those objects.

- Using Jacl:

```
$AdminTask createSSLConfig -interactive
```

- Using Jython:

```
AdminTask.createSSLConfig ('[-interactive]')
```

Example output:

Create a SSL Configuration.

```
*SSL Configuration Alias (alias): BIRKT40Node02SSLConfig
Management Scope Name (scopeName): (cell):BIRKT40Cell02:(node):BIRKT40Node02
Client Key Alias (clientKeyAlias): default
Server Key Alias (serverKeyAlias): default
SSL Type (type): [JSSE]
Client Authentication (clientAuthentication): [false]
Security Level of the SSL Configuration (securityLevel): [HIGH]
Enabled Ciphers SSL Configuration (enabledCiphers):
JSSE Provider (jsseProvider): [IBMJSSE2]
Client Authentication Support (clientAuthenticationSupported): [false]
SSL Protocol (sslProtocol): [SSL_TLS]
Trust Manager Object Names (trustManagerObjectNames): (cells/BIRKT40Cell02|security.xml#TrustManager_1)
*Trust Store Name (trustStoreName): NodeDefaultTrustStore
Trust Store Scope (trustStoreScopeName): (cell):BIRKT40Cell02:(node):BIRKT40Node02
*Key Store Name (keyStoreName): NodeDefaultKeyStore
Key Store Scope Name (keyStoreScopeName): (cell):BIRKT40Cell02:(node):BIRKT40Node02
Key Manager Name (keyManagerName): IbmX509
Key Manager Scope Name (keyManagerScopeName): (cell):BIRKT40Cell02:(node):BIRKT40Node02
```

Create SSL Configuration

```
F (Finish)
C (Cancel)
```

Select [F, C]: [F] F

```
WASX7278I: Generated command line: $AdminTask createSSLConfig {-alias BIRKT40Node02SSLConfig -scopeName
(cell):BIRKT40Cell02:(node):BIRKT40Node02 -clientKeyAlias default -serverKeyAlias default
-trustManagerObjectNames (cells/BIRKT40Cell02|security.xml#TrustManager_1) -trustStoreName
NodeDefaultTrustStore -trustStoreScopeName (cell):BIRKT40Cell02:(node):BIRKT40Node02 -keyStoreName
NodeDefaultKeyStore -keyStoreScopeName (cell):BIRKT40Cell02:(node):BIRKT40Node02 -keyManagerName
IbmX509 -keyManagerScopeName (cell):BIRKT40Cell02:(node):BIRKT40Node02 }
```

3. Save the configuration changes. See the Saving configuration changes with the wsadmin tool article for more information.
4. In a network deployment environment only, synchronize the node. See the Synchronizing nodes with the wsadmin tool article for more information.

Results

The name of the SSL configuration object that you created, for example, (cells/BIRKT40Cell02|security.xml#SSLConfig_1136652770753), appears in the security.xml file.

Example security.xml file output:

```
<repertoire xmi:id="SSLConfig_1136652770753" alias="BIRKT40Node02SSLConfig" type="JSSE"
managementScope="ManagementScope_1134610924357">
<setting xmi:id="SecureSocketLayer_1136652770924" clientKeyAlias="default" serverKeyAlias="default"
clientAuthentication="false" securityLevel="HIGH" jsseProvider="IBMJSSE2" sslProtocol="SSL_TLS"
keyStore="KeyStore_1134610924357" trustStore="KeyStore_1134610924377" trustManager="TrustManager_1"
keyManager="KeyManager_1134610924357"/>
</repertoire>
```

What to do next

Once you create the SSL configuration object, the next step is to use it. There are several different ways that you can associate SSL configurations with protocols, for example:

- Set the SSL configuration on the thread programmatically.
- Associate the SSL configuration with an outbound protocol or a target host and port.
- Directly associating the SSL configuration using the alias.
- Centrally managing the SSL configurations by associating them with SSL configuration groups or zones so that they are used based upon the group from where the end point exists.

Automating SSL configurations using scripting

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

Before you begin

Before starting this task, the wsadmin tool must be running. See the Starting the wsadmin scripting client article for more information.

About this task

AdminTask can be used in a interactive mode and batch mode. For automation the batch mode options should be used. AdminTask batch mode can be called in a JACL or Python script. Interactive mode will step through all the parameter the task needs, requires ones are marked with a '*'. Before the interactive task executes the task it echoes the batch mode syntax of the task to the screen. This can be helpful when writing batch mode scripts.

There attributes needed to create an ssl configurations:

- A key store
- Default client certificate alias
- Default server certificate alias
- Trust store
- The handshake protocol
- The ciphers needed during handshake
- Supporting client authentication or not

If automating the creation of a SSL Configuration it may be needed to create some of the attribute values needed like the key store, trust store, key manager, and trust managers.

- To create a SSL configuration the createSSLConfig AdminTask can be used. To make changes to the SSL configurations use the modifySSLConfig AdminTask.
 - Interactive mode:

Interactive mode steps you through all attributes and tell you the default value of the attribute if there is one. The default value is in '[' on the prompt line. The actual flag used in batch mode is in ')' on each prompt line. If you are using the default value then the flag will not show up on the batch command line.

Using Jacl:

```
$AdminTask createSSLConfig -interactive
```

– Using Jython:

```
AdminTask.createSSLConfig ('[interactive]')
```

Example output:

```
*SSL Configuration Alias (alias): testSSLConfig
Management Scope Name (scopeName): (cell):HOSTNode01Cell:(node):HOSTNode01
Client Key Alias (clientKeyAlias): clientCert
Server Key Alias (serverKeyAlias): serverCert
SSL Type (type): [JSSE]
Client Authentication (clientAuthentication): [false]
Security Level of the SSL Configuration (securityLevel): [HIGH] HIGH
Enabled Ciphers SSL Configuration (enabledCiphers):
JSSE Provider (jsseProvider): [IBMJSSE2]
Client Authentication Support (clientAuthenticationSupported): [false]
SSL Protocol (sslProtocol): [SSL_TLS] SSL_TLS
Trust Manager Object Names (trustManagerObjectNames):
*Trust Store Name (trustStoreName): testTrustStore
Trust Store Scope (trustStoreScopeName): (cell):HOSTNode01Cell:(node):HOSTNode01
*Key Store Name (keyStoreName): testKeyStore
Key Store Scope Name (keyStoreScopeName): (cell):HOSTNode01Cell:(node):HOSTNode01
Key Manager Name (keyManagerName): IbmX509
Key Manager Scope Name (keyManagerScopeName): (cell):HOSTNode01Cell:(node):HOSTNode01
```

Create SSL Configuration

F (Finish)

C (Cancel)

Select [F, C]: [F]

```
WASX7278I: Generated command line: $AdminTask createSSLConfig {-alias testSSLConfig
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01 -clientKeyAlias clientCert
-serverKeyAlias serverCert -trustStoreName testTrustStore
-trustStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyStoreName testKeyStore -keyStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyManagerName IbmX509 -keyManagerScopeName (cell):HOSTNode01Cell:(node):HOSTNode01 }
(cells/HOSTNode01Cell|security.xml#SSLConfig_1137687301834)
```

At the end of the output, the batch mode parameters are provided.

– Batch mode:

Using Jacl:

```
$AdminTask createSSLConfig {-alias testSSLConfig
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01 -clientKeyAlias clientCert
-serverKeyAlias serverCert -trustStoreName testTrustStore
-trustStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyStoreName testKeyStore -keyStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyManagerName IbmX509 -keyManagerScopeName (cell):HOSTNode01Cell:(node):HOSTNode01}
```

– Using Jython:

```
AdminTask.createSSLConfig ('[-alias testSSLConfig
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01 -clientKeyAlias clientCert
-serverKeyAlias serverCert -trustStoreName testTrustStore
-trustStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyStoreName testKeyStore -keyStoreScopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-keyManagerName IbmX509 -keyManagerScopeName (cell):HOSTNode01Cell:(node):HOSTNode01]')
```

Example output:

```
(cells/HOSTNode01Cell|security.xml#SSLConfig_1137687301834)
```

- Key Stores and Trust Stores The key store and trust store may already exist or a new one may need to be created. To create a new key store or trust store use the createKeyStore AdminTask. It will create a key store file and store the configuration object in the system configuration. A trust store is just a key store that usually only has signer certificates in it. To create a key store enter:

- Using Jacl:

```
$AdminTask createKeyStore {-keyStoreName testKeyStore -keyStoreType PKCS12
-keyStoreLocation $(USER_INSTALL_ROOT)\testKeyStore.p12 -keyStorePassword abcd
-keyStorePasswordVerify abcd -keyStoreIsFileBased true -keyStoreReadOnly false}
```

- Using Jython:

```
AdminTask.createKeyStore ('[-keyStoreName testKeyStore -keyStoreType PKCS12
-keyStoreLocation $(USER_INSTALL_ROOT)\testKeyStore.p12 -keyStorePassword abcd
-keyStorePasswordVerify abcd -keyStoreIsFileBased true -keyStoreReadOnly false]')
```

To populate the key store with certificates see “Managing Certificates using AdminConsole and Admin Task” The key store and trust store are required to create a SSL configuration. Use the '-keyStoreName' and '-trustStoreName' flags on the createSSLConfig. There scopes can be added with the '-keyStoreScope' flag and '-trustStoreScope' flags.

- Key Manager Key manager are used to determine how a certificate is selected. The IbmX509 key manager is in the security configuration by default. If a different key manager is needed then use createKeyManager AdminTask to create it. To create a key manager enter:

- Using Jacl:

```
$AdminTask createKeyManager {-name testKeyManager
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-provider IBMJSSE2 -algorithm specialAlgorithm }
```

- Using Jython:

```
AdminTask.createKeyManager ('[-name testKeyManager
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-provider IBMJSSE2 -algorithm specialAlgorithm]')
```

To supply a key manager on the createSSLConfig AdminTask use the '-keyManagerName' along with the '-keyManagerScope' flag.

- Trust Manager Trust managers are use to determine how trust is established during ssl communication. The IbmX509 and IbmPKIX are trust managers are in the security configuration by default. If a different or additional trust manager is needed then use the createTrustManger AdminTask to create it. To create a trust manager enter:

- Using Jacl:

```
$AdminTask createTrustManager {-name testTrustManager
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-provider IBMJSSE2 -algorithm specialAlgorithm }
```

- Using Jython:

```
AdminTask.createTrustManager ('[-name testTrustManager
-scopeName (cell):HOSTNode01Cell:(node):HOSTNode01
-provider IBMJSSE2 -algorithm specialAlgorithm]')
```

The SSL Configuration can have multiple trust managers. To supply multiple trust managers give a comma separated list of the trust managers configuration IDs with the -trustManagerObjectNames flag. When you create a trust manager the configuration object ID is returned. To get a list of trust managers object IDs use the **listTrustManagers** command of the AdminTask object with the -displayObjectName true flag. For example:

```
wsadmin>$AdminTask listTrustManagers -interactive
List Trust Managers
```

```
List trust managers.
```

```
Management Scope Name (scopeName):
Display list in ObjectName Format (displayObjectName): [false] true
```

```
List Trust Managers
```


F (Finish)
C (Cancel)

Select [F, C]: [F]

Inside generate script command

```
WASX7278I: Generated command line: $AdminTask listTrustManagers {-displayObjectName true }  
IbmX509(cells/IBM-0AF8DABCF16Node01Cell|security.xml#TrustManager_IBM-0AF8DABCF16Node01_1)  
IbmPKIX(cells/IBM-0AF8DABCF16Node01Cell|security.xml#TrustManager_IBM-0AF8DABCF16Node01_2)
```

Updating default key store passwords using scripting

Use the Jython or Jacl scripting language to change the default key store passwords. A key store file is created with a default password when you install the application server. Change this password to protect your security configuration.

Before you begin

Before starting this task, the wsadmin tool must be running. See the Starting the wsadmin scripting client article for more information.

About this task

When you install the application server, each server creates a key store and trust store for the default SSL configuration with the default password WebAS. To protect the security of the key store files and the SSL configuration, you must change the password. The following examples update the default password:

- Change multiple key stores passwords. The **changeMultipleKeyStorePasswords** command updates all of the key stores that have the same password. For example:

– Using Jacl:

```
$AdminTask changeMultipleKeyStorePasswords {-keyStorePassword WebAS  
-newKeyStorePassword secretPwd -newKeyStorePasswordVerify secretPwd}
```

– Using Jython:

```
AdminTask.changeMultipleKeyStorePasswords ['-keyStorePassword WebAS  
-newKeyStorePassword secretPwd -newKeyStorePasswordVerify secretPwd']
```

- Change the password of a single key store. The **changeKeyStorePassword** command updates the password of an individual key store. For example:

– Using Jacl:

```
$AdminTask changeKeyStorePassword {-keyStoreName testKS  
-keyStoreScope (cell):localhost:(server):server1  
-keyStorePassword WebAS -newKeyStorePassword secretPwd  
-newKeyStorePasswordVerify secretPwd}
```

– Using Jython:

```
AdminTask.changeKeyStorePassword ('[-keyStoreName testKS  
-keyStoreScope (cell):localhost:(server):server1  
-keyStorePassword WebAS -newKeyStorePassword secretPwd  
-newKeyStorePasswordVerify secretPwd]')
```

Configuring certificate authority client objects using the wsadmin tool

Use this topic to create a certificate authority (CA) client object. The client object contains all of the configuration information necessary to connect to your third-party CA server. A CA client must exist in your configuration before you can issue a request to the CA to create personal certificates with the requestCACertificate command.

Before you begin

A CA client object contains information that the system uses to connect to a certificate authority. Implement the com.ibm.ws.WSPKIClient interface to connect to the certificate authority and provide the

com.ibm.ws.WSPKIClient class when creating the CA client object.

About this task

If a CA client does not exist in your configuration, use the steps in this topic to create a new CA client.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine if a CA client exists in your configuration.

Use the following listCAClients command to list all certificate authority clients in your configuration:

```
print AdminTask.listCAClients()
```

3. If no CA clients exist, then create a new CA client.

Use the createCAClient command to create a new CA client object. The application server connects to a CA server through the WSPKIClient() implementation, which handles all connections and communications with the CA server. You must specify the following configuration information for a new CA client object:

Parameter	Description	Data Type
-CAClientName	Specify a name to uniquely identify the CA client object.	String

You can specify additional configuration information using the following parameters:

Parameter	Description	Data Type
-scopeName	Specify the management scope of the CA client. For a deployment manager profile, the system uses the cell scope as the default value. For an application server profile, the system uses the node scope as the default value.	String
-pkiClientImplClass	Specify the class path that implements the WSPKIClient interface. The system uses this path to connect to the CA and to issue requests to the CA. The default value is com.ibm.wsspi.ssl.WSPKIClient.	String
-host	Specify the host name in your system where the CA resides.	String
-port	Specify the port on the server where the CA listens.	String
-userName	Specify the user name to use to authenticate to the CA.	String
-password	Specify the password for the user name that authenticates to the CA.	String
-frequencyCheck	Specify how often, in minutes, the system checks with the CA to determine if a certificate has been created.	String
-retryCheck	Specify the number of times to check with the CA to determine if a certificate has been created.	String
-customProperties	Specifies a comma separated list of attribute and value (attribute=value) custom property pairs to add to the CA client object.	String

Use the following example command to create a new CA client object:

```
AdminTask.createCAClient('[-caClientName clientObj01 -pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient -host machine011 -port 9022 -userName admin -password pw4admin]')
```

The command returns the object name of the CA client that has been created.

4. Save your configuration changes.

What to do next

If the CA client object was successfully created, then you can configure the application server to use a personal certificate created by an external CA.

Administering certificate authority clients using the wsadmin tool

Use this topic to modify certificate authority (CA) client objects. The client object contains all of the configuration information necessary to connect to your third-party CA server.

Before you begin

You must configure a CA client object in your environment.

About this task

For existing CA client objects, use the steps in this topic to view, modify, or delete existing CA client object configurations.

- View existing CA client objects and configuration data. Use the `listCAClients` and `getCAClient` commands to query your environment for your existing CA clients.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. List all CA client objects in your configuration.

Use the `listCAClients` command to list all certificate authority clients in your configuration. If you do not provide a value for the `-scopeName` parameter, then the command queries the cell if you use a deployment manager profile or queries the node if you use an application server profile. Use the `-all` parameter to query your environment without using a specific scope, as the following example demonstrates:

```
print AdminTask.listCAClients('-all true')
```

The command returns an array of attribute lists, displaying one attribute list for each CA client, as the following example output displays:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell):myCell101] [name jenCAClient] [baseDn ] [_Websphere_Config_Dat
a_Id cells/myCell101|security.xml#CAClient_1181834566881] [port 2950] [CACertifi
cate ] [pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Webspher
e_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pa
ssword ] [host ] ]'
```

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell):myCell101] [name myCAClient] [baseDn ] [_Websphere_Config_Dat
a_Id cells/myCell101|security.xml#CAClient_1181834566882] [port 2951] [CACertific
ate ] [pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Websphere
_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pas
sword ] [host ] ]'
```

3. List the configuration attributes for a specific CA client.

Use the `getCAClient` command to view the list of attributes for a specific CA client, as the following example demonstrates:

```
print AdminTask.getCAClient('-caClientName myCAClient')
```

The command returns an attribute list that contains the attribute and value pairs for the specific CA client, as the following example demonstrates:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementSc
ope_1)] [scopeName (cell):myCell101] [name myCAClient] [baseDn ] [_Webspher
e_Config_Data_Id cells/myCell101|security.xml#CAClient_1181834566882] [por
t 2951] [CACertificate ] [pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [u
serId ] [_Websphere_Config_Data_Type CAClient] [retryCheck 0] [properties ] [fre
quencyCheck 0] [password ] [host ] ]'
```

- Modify your existing CA client object configuration data. Use the `modifyCAClient` command to change one or more configuration attributes for a specific CA client.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine which configuration attributes to edit.

The `modifyCAClient` modifies all attributes that you specify with the command parameters. If you do not specify a parameter, then its corresponding attribute does not change. You can edit the following configuration data with the `modifyCAClient` command:

Parameter	Description	Data Type
-scopeName	Specify the management scope of the CA client. For a deployment manager profile, the system uses the cell scope as the default. For an application server profile, the system uses the node scope as the default.	String
-pkiClientImplClass	Specify the class path that implements the <code>WSPKIClient</code> interface. The system uses this path to connect to the CA and to issue requests to the CA.	String
-host	Specify the host name in your system where the CA resides.	String
-port	Specify the port on the server where the CA listens.	String
-userName	Specify the user name to use to authenticate to the CA.	String
-password	Specify the password for the user name that authenticates to the CA.	String
-frequencyCheck	Specify how often, in minutes, the system should check with the CA to determine if a certificate has been created.	String
-retryCheck	Specify the number of times to check with the CA to determine if a certificate has been created.	String
-customProperties	Specifies a comma separated list of attribute and value (<code>attribute=value</code>) custom property pairs to modify on the CA Client object. You can create, modify, or remove properties. To remove a property specify <code>attribute= attribute</code> as equal to no value.	String

3. Modify specific configuration attributes for a CA client object.

Use the following example command to modify the port number of the CA, the user name, and password attributes for the `myCAClient` CA client object:

```
AdminTask.modifyCAClient('[-caClientName myCAClient -port 4060 -userName admin -password password4admin -pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient]')
```

4. Save your configuration changes.

- Remove a CA client object from your configuration. Use the `deleteCAClient` command to delete a CA client object from your configuration. The command does not delete the CA client object if the CA client to delete is referenced by a certificate object.

- Launch the `wsadmin` scripting tool using the Jython scripting language.
- Determine the CA client object to delete.

Use the `listCAClients` command to list all certificate authority clients in your configuration. If you do not provide a value for the `-scopeName` parameter, then the command queries the cell if you use a deployment manager profile or queries the node if you use an application server profile. Use the `-all` parameter to query your environment without using a specific scope, as the following example demonstrates:

```
print AdminTask.listCAClients('-all true')
```

3. Delete the CA client object of interest.

Use the `deleteCAClient` command to delete the CA client object from your configuration. Use the `-caClientName` parameter to specify the CA client to delete. You can optionally specify the management scope of the CA client object with the `scopeName` parameter. The following example command removes the `myCAClient` CA client object:

```
AdminTask.deleteCAClient('[-caClientName myCAClient]')
```

If you receive an error message, then verify that the CA client object of interest exists in your configuration and that it is not referenced by a certificate object in your security configuration.

4. Save your configuration changes.

Setting a certificate authority certificate as the default certificate using the wsadmin tool

Use this topic to make a request to an external certificate authority (CA) to create a personal certificate. After the CA returns the certificate and the certificate is saved in the keystore, then you can use it as the server default personal certificate.

Before you begin

You must configure a CA client object in your environment. The client object contains all of the configuration information necessary to connect to your third-party CA server.

About this task

After profile creation, the system is assigned a default chained personal certificate. Use the following steps to modify the application server to use a default personal certificate created by an external CA.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Verify that a certificate authority client exists in your configuration. Use the `listCAClients` command to query your environment for all existing certificate authority clients and configuration attributes, or the `getCAClient` command to return the configuration attributes for a specific certificate authority client. If the `listCAClients` or `getCAClient` commands do not return any attributes, then you must create a certificate authority client object before you can complete the remaining steps.

- List all certificate authority client objects in your configuration.

Use the `listCAClients` command to list all certificate authority clients in your configuration. If you do not provide a value for the `-scopeName` parameter, then the command queries the cell if you use a deployment manager profile or queries the node if you use an application server profile. Use the `-all` parameter to query your environment without using a specific scope, as the following example demonstrates:

```
print AdminTask.listCAClients('-all true')
```

The command returns an array of attribute lists, displaying one attribute list for each CA client, as the following example output displays:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell):myCell101] [name jenCAClient] [baseDn ] [_Websphere_Config_Data
Data_Id cells/myCell101|security.xml#CAClient_1181834566881] [port 2950] [CACertifi
cate ] [pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Webspher
e_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pa
ssword ] [host ] ]'
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell):myCell101] [name myCAClient] [baseDn ] [_Websphere_Config_Dat
a_Id cells/myCell101|security.xml#CAClient_1181834566882] [port 2951] [CACertific
ate ] [pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Websphere
_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pas
sword ] [host ] ]'
```

- List the configuration attributes for a specific certificate authority client.

Use the `getCAClient` command to view the list of attributes for a specific certificate authority client, as the following example demonstrates:

```
print AdminTask.getCAClient('-caClientName myCAClient')
```

The command returns an attribute list that contains the attribute and value pairs for the specific certificate authority client, as the following example demonstrates:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementSc
ope_1)] [scopeName (cell):myCell101] [name myCAClient] [baseDn ] [_Websphe
re_Config_Data_Id cells/myCell101|security.xml#CAClient_1181834566882] [por
t 2951] [CACertificate ] [pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [u
serId ] [_Websphere_Config_Data_Type CAClient] [retryCheck 0] [properties ] [fre
quencyCheck 0] [password ] [host ] ]'
```

3. Optional: If a certificate authority client does not exist in your environment, then configure a CA client object.
4. Optional: View the current default personal certificate.

Use the following `listPersonalCertificates` command to display the current default personal certificate to replace:

```
AdminTask.listPersonalCertificates('[-keyStoreName CellDefaultKeyStore -keyStoreScope (cell):myCell01]')
```

5. Request a certificate from a certificate authority.

Before the current default personal certificate can be replaced, you must request a certificate from a certificate authority. You can create a new certificate request or use the `createCertificateRequest` command to use a predefined certificate request. The system uses the certificate request and the certificate authority configuration information from the CA client object to request the certificate from the certificate authority. If the certificate authority returns a certificate, then the `requestCACertificate` command stores the certificate in the specified key store and returns a message of COMPLETE. Use the `requestCACertificate` command and the following required parameters to request a certificate from a certificate authority:

Parameter	Description	Data Type
-certificateAlias	Specifies the alias of the certificate. You can specify a predefined certificate request.	String
-keyStoreName	Specifies the name of the keystore object that stores the CA certificate. Use the <code>listKeyStores</code> command to display a list of available keystores.	String
-caClientName	Specifies the name of the CA client that was used to create the CA certificate.	String
-revocationPassword	Specifies the password to use to revoke the certificate at a later date.	String

You can also use the following parameters to specify additional certificate request options. If you do not specify an optional parameter, then the command uses the default value.

Parameter	Description	Data Type
-keyStoreScope	Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope.	String
-caClientScope	Specifies the management scope of the CA client. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope.	String
-certificateCommonName	Specifies the common name (CN) part of the full distinguished name (DN) of the certificate. This common name can represent a person, company, or machine. For Web sites, the common name is frequently the DNS host name where the server resides.	String
-certificateSize	Specifies the size of the certificate key. The valid values are 512, 1024, and 2048. The default value is 1024.	String
-certificateOrganization	Specifies the organization portion of the distinguished name.	String
-certificateOrganizationalUnit	Specifies the organizational unit portion of the distinguished name.	String
-certificateLocality	Specifies the locality portion of the distinguished name.	String
-certificateState	Specifies the state portion of the distinguished name.	String
-certificateZip	Specifies the zip code portion of the distinguished name.	String
-certificateCountry	Specifies the country portion of the distinguished name.	String

Use the following example command syntax to request a certificate from a certificate authority:

```
AdminTask.requestCACertificate('-certificateAlias newCertificate -keyStoreName CellDefaultKeyStore -caClientName myCAClient -revocationPassword revokeCApw -pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient')
```

The command returns one of two values: Certificate COMPLETE or certificate PENDING. If the command returns the Certificate COMPLETE message, the certificate authority returned the requested certificate and the default personal certificate is replaced. If the command returns the certificate PENDING message, the certificate authority did not yet return a certificate. Use the queryCACertificate command to view the current status of the certificate request, as the following example demonstrates:

```
AdminTask.queryCACertificate('-certificateAlias newCertificate -keyStoreName
CellDefaultKeyStore -pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient')
```

6. Replace the server default personal certificate.

Use the following replaceCertificate command example to replace the existing default personal certificate with the newly created CA personal certificate:

```
AdminTask.replaceCertificate('-keyStoreName CellDefaultKeyStore -certificateAlias
defaultPersonalCertificate -replacementCertificateAlias newCertificate')
```

7. Save your configuration changes.

Results

The default personal certificate for the server is a certificate that is created by an external CA.

What to do next

If the CA client object was successfully created, then you can configure the application server to use a personal certificate created by an external CA.

Creating certificate authority (CA) personal certificates using the wsadmin tool

Use this topic to create CA certificates from a certificate authority (CA).

Before you begin

You must configure a CA client object in your environment. The client object contains all of the configuration information necessary to connect to your third-party CA server.

About this task

Use the following information to create a CA personal certificate using a CA client.

1. Optional: Query your configuration for keystores to determine where system stores the new CA certificate.

Use the listKeyStores command to list all keystores for a specific management scope. Specify the -scopeName parameter to display keystores within a specific management scope, or set the -all parameter to true to display all keystores regardless of scope. The following example lists all keystores in your configuration:

```
AdminTask.listKeyStores('-all true')
```

The command returns the following sample output:

```
CellDefaultKeyStore(cells/myCell|security.xml#KeyStore_1)
CellDefaultTrustStore(cells/myCell|security.xml#KeyStore_2)
CellLTPAKeys(cells/myCell|security.xml#KeyStore_3)
NodeDefaultKeyStore(cells/myCell|security.xml#KeyStore_1598745926544)
NodeDefaultTrustStore(cells/myCell|security.xml#KeyStore_1476529854789)
```

Use the getKeyStoreInfo command and specify the -keyStoreName parameter to return additional information about the keystore of interest, as the following example displays:

```
AdminTask.getKeyStoreInfo('[-keyStoreName CellDefaultKeyStore]')
```

The command returns the following configuration information for the keystore of interest:

```
[ [location ${CONFIG_ROOT}/cells/myCell/key.p12] [password *****] [_Websphere_Config_Data_Id cells/myCell|security.xml#KeyStore_1] [_Websphere_Config_Data_Version ] [useForAcceleration false] [slot 0] [type PKCS12] [additionalKeyStoreAttrs ] [fileBased true] [_Websphere_Config_Data_Type KeyStore] [customProviderClass ] [hostList ] [createStashFileForCMS false] [description [Default key store for JenbCell101]] [readOnly false] [initializeAtStartup false] [managementScope (cells/JenbCell101|security.xml#ManagementScope_1)] [usage SSLKeys] [provider IBMJCE] [name CellDefaultKeyStore] ]
```

2. Optional: Determine which CA client to use.

Use the `listCAClients` command to list the CA clients that exist in your configuration. Specify the `-scopeName` parameter to display CA clients within a specific management scope, or set the `-all` parameter to `true` to display all CA clients regardless of scope. The following example lists all CA clients in your configuration:

```
AdminTask.listCAClients('-all true')
```

3. Create a CA personal certificate.

Use the `requestCACertificate` command to create a new CA personal certificate in your environment. The system uses the certificate request and the certificate authority configuration information from the CA client object to request the certificate from the certificate authority. If the certificate authority returns a certificate, the `requestCACertificate` command stores the certificate in the specified key store and returns a message of `COMPLETE`. Use the `requestCACertificate` command and the following required parameters to request a certificate from a certificate authority:

Parameter	Description	Data type
<code>-certificateAlias</code>	Specifies the alias of the certificate. You can specify a predefined certificate request.	String
<code>-keyStoreName</code>	Specifies the name of the keystore object that stores the CA certificate. Use the <code>listKeyStores</code> command to display a list of available keystores.	String
<code>-caClientName</code>	Specifies the name of the CA client that was used to create the CA certificate.	String
<code>-revocationPassword</code>	Specifies the password to use to revoke the certificate at a later date.	String

You can also use the following parameters to specify additional certificate request options. If you do not specify an optional parameter, the command uses the default value.

Parameter	Description	Data type
<code>-keyStoreScope</code>	Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope.	String
<code>-caClientScope</code>	Specifies the management scope of the CA client. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope.	String
<code>-certificateCommonName</code>	Specifies the common name (CN) part of the full distinguished name (DN) of the certificate. This common name can represent a person, company, or machine. For Web sites, the common name is frequently the DNS host name where the server resides.	String
<code>-certificateSize</code>	Specifies the size of the certificate key. The valid values are 512, 1024, and 2048. The default value is 1024.	String
<code>-certificateOrganization</code>	Specifies the organization portion of the distinguished name.	String
<code>-certificateOrganizationalUnit</code>	Specifies the organizational unit portion of the distinguished name.	String
<code>-certificateLocality</code>	Specifies the locality portion of the distinguished name.	String
<code>-certificateState</code>	Specifies the state portion of the distinguished name.	String

Parameter	Description	Data type
-certificateZip	Specifies the zip code portion of the distinguished name.	String
-certificateCountry	Specifies the country portion of the distinguished name.	String

Use the following example command syntax to request a certificate from a certificate authority:

```
AdminTask.requestCACertificate('-certificateAlias newCertificate -keyStoreName CellDefaultKeyStore -CAClientName myCAClient -revocationPassword revokeCApw')
```

The command returns one of two values: Certificate COMPLETE or certificate PENDING. If the command returns the Certificate COMPLETE message, the certificate authority returned the requested certificate and the default personal certificate is replaced. If the command returns the certificate PENDING message, the certificate authority did not yet return a certificate. Use the queryCACertificate command to view the current status of the certificate request, as the following example displays:

```
AdminTask.queryCACertificate('-certificateAlias newCertificate -keyStoreName CellDefaultKeyStore')
```

4. Save your configuration changes.

Results

The default personal certificate for the server is a certificate that is created by an external CA.

Revoking certificate authority personal certificates using the wsadmin tool

You can revoke CA certificates from a certificate authority (CA). Revoke personal certificates that are no longer being used in your configuration.

Before you begin

Use the requestCACertificate command to create a personal certificate with the requestCACertificate task before you can request that the certificate authority revoke the certificate. Certificates created with the requestCACertificate command have an associated reference object in the configuration that you can use to submit the certificate revocation request to the certificate authority.

About this task

This topic uses the revokeCACertificate command to submit a request to revoke a certificate on the certificate authority. You can only revoke a certificate that was created with the requestCACertificate command. You must specify the revocation password that was provided when the certificate was created. Use the same password to revoke the certificate on the certificate authority.

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Determine the CA personal certificate to revoke.

Use the listPersonalCertificates command to view a list of all personal certificates and associated attributes for a specific keystore, as the following example demonstrates:

```
AdminTask.listPersonalCertificates('-keyStoreName CellDefaultKeyStore')
```

The command returns an attribute list for each personal certificate, including CA personal certificates. CA personal certificates only return the status attribute. You can revoke each CA personal certificates that returns a COMPLETE status. Determine which CA personal certificate to revoke.

3. Revoke a CA personal certificate.

Use the revokeCACertificate command to revoke the CA personal certificate of interest. You must specify the name of the keystore, certificate alias, and revocation password using the following parameters:

Parameter	Description	Data Type
-keyStoreName	Specifies the name of the keystore where the CA personal certificate is stored.	String

Parameter	Description	Data Type
-certificateAlias	Specifies the unique name that identifies the CA personal certificate object and the alias name of the certificate in the keystore.	String
-revocationPassword	Specifies the password needed to revoke the certificate. This is the same password that was provided when the certificate was created.	String

You can specify additional information with the following optional parameters:

Parameter	Description	Data Type
-keyStoreScope	Specifies the management scope of the keystore. For a deployment manager profile, the system uses the cell scope as the default value. For an application server profile, the system uses the node scope as the default value.	String
-revocationReason	Specifies the reason for revoking the certificate of interest. The default value for this parameter is unspecified.	String

The following example revokes a CA personal certificate:

```
AdminTask.revokeCACertificate(['-keyStoreName CellDefaultKeyStore -certificateAlias myCertificate -revocationPassword pw4revoke'])
```

4. Save your configuration changes.

CAClientCommands command group for the AdminTask object

You can use the Jython scripting language to manage your certificate authority (CA) client configurations with the wsadmin tool. Use the commands and parameters in the CAClientCommands group to create, modify, query, and remove connections to a third-party CA server.

Use the following commands to manage your certificate authority (CA) client configurations:

- “createCAClient”
- “modifyCAClient” on page 973
- “getCAClient” on page 974
- “deleteCAClient” on page 975
- “listCAClients” on page 976

createCAClient

The createCAClient command creates a new CA client object in your configuration. The application server connects to a CA server through the WSPKIClient() implementation, which handles all connections and communications with the CA server.

Target object

None.

Required parameters

-caClientName

Specifies a name to uniquely identify the CA client object. (String, required)

-pkiClientImplClass

Specifies the class path that implements the WSPKIClient interface. The system uses this path to connect to the CA and to issue requests to the CA. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the CA client. For a deployment manager profile, the system uses the cell scope as the default value. For an application server profile, the system uses the node scope as the default value. (String, optional)

-host

Specifies the host name in your system where the CA resides. (String, optional)

-port

Specifies the port on the server where the CA listens. (String, optional)

-userName

Specifies the user name to use to authenticate to the CA. (String, optional)

-password

Specifies the password for the user name that authenticates to the CA. (String, optional)

-frequencyCheck

Specifies how often, in minutes, the system communicates with the CA to determine if a certificate has been created. (String, optional)

-retryCheck

Specifies the number of times to communicate with the CA to determine if a certificate has been created. (String, optional)

-customProperties

Specifies a comma-separated list of attribute and value custom property pairs to add to the CA client object, using the following format: attribute=value,attribute=value. (String, optional)

Return value

The command returns the object name of the CA client that the system creates.

Batch mode example usage

- Using Jython string:

```
AdminTask.createCAClient(['-caClientName clientObj01 -pkiClientImplClass
com.ibm.wsspi.ssl.WSPKIClient -host machine011 -port 9022
-userName admin -password pw4admin'])
```

- Using Jython list:

```
AdminTask.createCAClient(['-caClientName', 'clientObj01', '-pkiClientImplClass',
'com.ibm.wsspi.ssl.WSPKIClient', '-host', 'machine011', '-port', '9022',
-userName', 'admin', '-password', 'pw4admin'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createCAClient('-interactive')
```

modifyCAClient

The modifyCAClient command modifies your existing CA client object configuration data. You can modify one or multiple configuration attributes for a specific CA client.

Target object

None.

Required parameters

-caClientName

Specifies the name of the CA client of interest. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the CA client. For a deployment manager profile, the system uses the cell scope as the default. For an application server profile, the system uses the node scope as the default. (String, optional)

-pkiClientImplClass

Specifies the class path that implements the WSPKIClient interface. The system uses this path to connect to the CA and to issue requests to the CA. (String, optional)

-host

Specifies the host name in your system where the CA resides. (String, optional)

-port

Specifies the port on the server where the CA listens. (String, optional)

-userName

Specifies the user name to use to authenticate to the CA. (String, optional)

-password

Specifies the password for the user name that authenticates to the CA. (String, optional)

-frequencyCheck

Specifies how often, in minutes, the system should check with the CA to determine if a certificate has been created. (String, optional)

-retryCheck

Specifies the number of times to check with the CA to determine if a certificate has been created. (String, optional)

-customProperties

Specifies a comma separated list of attribute and value (attribute=value) custom property pairs to modify on the CA Client object. You can create, modify, or remove properties. To remove a property specify the attribute and value as attribute=. (String, optional)

Return value

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyCAClient(['-caClientName myCAClient -port 4060  
-userName admin -password password4admin'])
```

- Using Jython list:

```
AdminTask.modifyCAClient(['-caClientName', 'myCAClient', '-port', '4060',  
'-userName', 'admin', '-password', 'password4admin'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyCAClient('-interactive')
```

getCAClient

The getCAClient command displays a list of attributes for a specific CA client.

Target object

None.

Required parameters

-caClientName

Specifies the CA client name of interest. (String, required)

Optional parameters

-scopeName

Specifies the management scope of CA client of interest. (String, optional)

Return value

The command returns an attribute list that contains the attribute and value pairs for the specific CA client, as the following example displays:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)] [scopeName (cell):myCell101] [name myCAClient] [baseDn ] [WebSphere_Config_Data_Id cells/myCell101|security.xml#CAClient_1181834566882] [port 2951] [CACertificate ] [pkIClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [WebSphere_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [password ] [host ] ]'
```

Batch mode example usage

- Using Jython string:

```
print AdminTask.getCAClient('-caClientName myCAClient')
```

- Using Jython list:

```
print AdminTask.getCAClient('-caClientName', 'myCAClient')
```

Interactive mode example usage

- Using Jython string:

```
print AdminTask.getCAClient('-interactive')
```

deleteCAClient

The deleteCAClient command removes the CA client object of interest from your configuration. Use the -caClientName parameter to specify the CA client to delete. You can optionally specify the management scope of the CA client object with the scopeName parameter.

Target object

None.

Required parameters

-caClientName

Specifies the name of the CA client of interest. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the CA client of interest. (String, optional)

Return value

The command does not return output if the system successfully removes the CA client of interest. If you receive an error message, verify that the CA client object of interest exists in your configuration and that it is not referenced by a certificate object in your security configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteCAClient('[-caClientName myCAClient]')
```

- Using Jython list:

```
AdminTask.deleteCAClient(['-caClientName', 'myCAClient'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteCAClient('-interactive')
```

listCAClients

The listCAClients command lists all CA clients in your configuration or within a specific scope. If you do not provide a value for the -scopeName parameter, the command queries the cell if you use a deployment manager profile or queries the node if you use an application server profile. Use the -all parameter to query your environment without using a specific scope.

Target object

None.

Optional parameters

-scopeName

Specifies the management scope to search for CA clients. (String, optional)

-all

Specifies whether the system queries for CA clients without a specific scope. (Boolean, optional)

Return value

The command returns an array of attribute lists, displaying one attribute list for each CA client, as the following example output displays:

```
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell1):myCell101] [name jenCAClient] [baseDn ] [_Websphere_Config_Da
ta_Id cells/myCell101|security.xml#CAClient_1181834566881] [port 2950] [CACertifi
cate ] [pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Webspher
e_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pa
ssword ] [host ] ]'
'[ [backupCAs ] [managementScope (cells/myCell101|security.xml#ManagementScope_1)
] [scopeName (cell1):myCell101] [name myCAClient] [baseDn ] [_Websphere_Config_Dat
a_Id cells/myCell101|security.xml#CAClient_1181834566882] [port 2951] [CACertific
ate ] [pkiClientImplClass com.ibm.wsspi.ssl.WSPKIClient] [userId ] [_Websphere
_Config_Data_Type CAClient] [retryCheck 0] [properties ] [frequencyCheck 0] [pas
sword ] [host ] ]'
```

Batch mode example usage

- Using Jython string:

```
print AdminTask.listCAClients('-all true')
```

- Using Jython list:

```
print AdminTask.listCAClients('-all', 'true')
```

Interactive mode example usage

- Using Jython:

```
print AdminTask.listCAClients('-interactive')
```

Creating self-signed certificates using scripting

Use the Jython or Jacl scripting language to create self-signed certificates with the wsadmin tool.

Before you begin

Before starting this task, the wsadmin tool must be running. See the Starting the wsadmin scripting client article for more information.

About this task

You can create self-signed certificates using the scripting and the AdminTask object. You can run the commands in interactive or batch mode. Interactive mode provides a way to discover the flags that you need to run the task in batch mode.

Certificates reside inside of key stores. To run the commands, you will need the name of the key store to be supplied. Use the **listKeyStore** command of the AdminTask object to get a list of key stores. If you need a new key store, use the **createKeyStore** command of the AdminTask object.

To create a personal key store, use the following examples:

- Interactive mode:

- Using Jython:

```
AdminTask.createSelfSignedCertificate ('[-interactive]')
```

- Using Jacl:

```
$AdminTask createSelfSignedCertificate -interactive
```

Example output:

```
*Key Store Name (keyStoreName): keyStore
Key Store Scope Name (keyStoreScope):
*Certificate Alias (certificateAlias): newCert
"Certificate Version" (certificateVersion): 3
*Key Size (certificateSize): [1024]
*Common Name (certificateCommonName): localhost
*Organization (certificateOrganization): workgroup
Organizational Unit (certificateOrganizationalUnit): testing
certLocality (certificateLocality): austin
State (certificateState): Texas
Zip (certificateZip): 78757
Country (certificateCountry): [US]
Validity Period (certificateValidDays): [365]
Create Self-Signed Certificate
```

```
F (Finish)
```

```
C (Cancel)
```

```
Select [F, C]: [F]
```

```
WASX7278I: Generated command line: $AdminTask createSelfSignedCertificate
{-keyStoreName keyStore -certificateAlias newCert -certificateVersion 3
-certificateCommonName localhost -certificateOrganization ibm
-certificateOrganizationalUnit testing -certificateLocality austin
-certificateState Texas -certificateZip 78757 }
true
```

At the end of the output, the batch mode parameters are provided.

- Batch mode:

- Using Jython:

```
AdminTask.createSelfSignedCertificate ('[-keyStoreName keyStore
-certificateAlias newCert -certificateVersion 3 -certificateSize 1024
-certificateCommonName localhost -certificateOrganization ibm
-certificateOrganizationalUnit testing -certificateLocality austin
-certificateState Texas -certificateZip 78757]')
```

- Using Jacl:

```
$AdminTask createSelfSignedCertificate {-keyStoreName keyStore
-certificateAlias newCert -certificateVersion 3 -certificateSize 1024
-certificateCommonName localhost -certificateOrganization ibm
-certificateOrganizationalUnit testing -certificateLocality austin
-certificateState Texas -certificateZip 78757 }
```


keyManagerCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security. The commands and parameters in the keyManagerCommands group can be used to manage key manager settings. You can use these commands to create, modify, list, or obtain information about key managers.

The keyManagerCommands command group for the AdminTask object includes the following commands:

- “deleteKeyManager”
- “getKeyManager”
- “listKeyManagers” on page 979
- “modifyKeyManager” on page 980

deleteKeyManager

The **deleteKeyManager** command deletes the key manager settings from the configuration.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key manager. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyManager {-name testKM}
```

- Using Jython string:

```
AdminTask.deleteKeyManager(['-name testKM'])
```

- Using Jython list:

```
AdminTask.deleteKeyManager(['-name', 'testKM'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyManager {-interactive}
```

- Using Jython:

```
AdminTask.deleteKeyManager('-interactive')
```

getKeyManager

The **getKeyManager** command displays a properties object that contains the key manager attributes and values.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key manager. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getKeyManager {-name testKM}
```

- Using Jython string:

```
AdminTask.getKeyManager(['-name testKM'])
```

- Using Jython list:

```
AdminTask.getKeyManager(['-name', 'testKM'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getKeyManager {-interactive}
```

- Using Jython:

```
AdminTask.getKeyManager('-interactive')
```

listKeyManagers

The **listKeyManagers** command lists the key managers within a particular management scope.

Target object

None.

Required parameters

None.

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-displayObjectName

Set the value of this parameter to `true` to list the key manager objects within the scope. Set the value of this parameter to `false` to list the strings that contain the key manager name and the management scope. (Boolean, optional)

-all

Specify the value of this parameter as `true` to list all key managers. This parameter overrides the `scopeName` parameter. The default value is `false`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listKeyManagers
```

- Using Jython:

```
AdminTask.listKeyManagers()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listKeyManagers {-interactive}
```

- Using Jython:

```
AdminTask.listKeyManagers('-interactive')
```

modifyKeyManager

The **modifyKeyManager** command changes existing key manager settings.

Target object

None.

Required parameters

-name

The name that uniquely identifies the key manager. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-provider

Specifies the provider name of the key manager. (String, optional)

-algorithm

Specifies the algorithm name of the key manager. (String, optional)

-keyManagerClass

Specifies the name of the key manager implementation class. You cannot use this parameter with the provider or the algorithm parameter. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifyKeyManager {-name testKM -provider IBMJSSE2 -algorithm IbmX509}
```

- Using Jython string:

```
AdminTask.modifyKeyManager(['-name testKM -provider IBMJSSE2 -algorithm IbmX509'])
```

- Using Jython list:

```
AdminTask.modifyKeyManager(['-name', 'testKM', '-provider', 'IBMJSSE2', '-algorithm', 'IbmX509'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifyKeyManager {-interactive}
```

- Using Jython:

```
AdminTask.modifyKeyManager('-interactive')
```

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

“Creating an SSL configuration at the node scope using scripting” on page 957

An Secure Socket Layer (SSL) configuration references many other configuration objects. To help you make valid selections for the new SSL configuration before you create it, view information about existing configuration objects. Information about existing objects is also useful when you create a node scoped SSL configuration using the **createSSLConfig** command of the AdminTask object.

KeyStoreCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure keystores with the wsadmin tool. A keystore is created by the application server during install and can contain cryptographic keys or certificates. The commands and parameters in the KeyStoreCommands group can be used to create, delete, and manage keystores.

The KeyStoreCommands command group for the AdminTask object includes the following commands:

- “changeKeyStorePassword”
- “changeMultipleKeyStorePasswords” on page 982
- “createKeyStore” on page 983
- “createCMSKeyStore” on page 984
- “deleteKeyStore” on page 985
- “exchangeSigners” on page 986
- “getKeyStoreInfo” on page 986
- “listKeyFileAliases” on page 987
- “listKeyStores” on page 988
- “listKeyStoreTypes” on page 988
- “modifyKeyStore” on page 989

changeKeyStorePassword

The **changeKeyStorePassword** command modifies the password of a keystore. The command automatically saves the new password to the configuration.

Required parameters

-keyStoreName

Specifies the name of the password to change. (String, required)

-keyStorePassword

Specifies the name of the password to change. (String, required)

-newKeyStorePassword

Specifies the new password that to use to access the keystore. (String, required)

-newKeyStorePasswordVerify

Specifies the new password to confirm the new keystore password. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the keystore. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask changeKeyStorePassword {-keystoreName myKeystore -keyStorePassword
WebAS -newKeyStorePassword newpwd -newKeyStorePasswordVerify newpwd}
```

- Using Jython string:

```
AdminTask.changeKeyStorePassword(['-keystoreName myKeystore -keyStorePassword
WebAS -newKeyStorePassword newpwd -newKeyStorePasswordVerify newpwd'])
```

- Using Jython list:

```
AdminTask.changeKeyStorePassword(['-keystoreName', 'myKeystore', '-keyStorePassword',
'WebAS', '-newKeyStorePassword', 'newpwd', '-newKeyStorePasswordVerify', 'newpwd'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask changeKeyStorePassword {-interactive}
```

- Using Jython:

```
AdminTask.changeKeyStorePassword('-interactive')
```

changeMultipleKeyStorePasswords

The **changeMultipleKeyStorePasswords** command updates the passwords for each keystores in the configuration that has a specific password. This is useful because when you create keystore files on the system, they will have WebAS as a password by default.

Required parameters

-keyStorePassword

Specifies the name of the password that you want to change. (String, required)

-newKeyStorePassword

Specifies the new password that you will use to access the keystore. (String, required)

-newKeyStorePasswordVerify

Confirms the new keystore password. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask changeMultipleKeyStorePasswords {-keyStorePassword WebAS
-newKeyStorePassword newpwd -newKeyStorePasswordVerify newpwd}
```

- Using Jython string:

```
AdminTask.changeMultipleKeyStorePasswords(['-keyStorePassword WebAS
-newKeyStorePassword newpwd -newKeyStorePasswordVerify newpwd'])
```

- Using Jython list:

```
AdminTask.changeMultipleKeyStorePasswords(['-keyStorePassword', 'WebAS',
'-newKeyStorePassword', 'newpwd', '-newKeyStorePasswordVerify', 'newpwd'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask changeMultipleKeyStorePasswords {-interactive}
```

- Using Jython:

```
AdminTask.changeMultipleKeyStorePasswords('-interactive')
```

createKeyStore

The **createKeyStore** command creates the keystore settings in the configuration and the keystore database.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-keyStoreType

The implementation of the keystore management. (String, required)

-keyStoreLocation

The location of the keystore. For file based, the location is the files system path to the keystore database. For hardware keystore, the location is the path to the token library. (String, required)

-keyStorePassword

The password that protects the keystore. (String, required)

-keyStorePasswordVerify

The password that protects the keystore. (String, required)

Optional parameters

-keyStoreProvider

The provider used to implement the keystore. (String, optional)

-keyStoreIsFileBased

Set the value of this parameter to `true` if the keystore is file based. Set the value of this parameter to `false` for hardware crypto keystores. (Boolean, optional)

-keyStoreHostList

A list of host names that indicate from where the keystore is remotely managed, separated by commas. (String, optional)

-keyStoreInitAtStartup

Set the value of this parameter to `true` if the keystore is initialized at startup. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

-keyStoreReadOnly

Set the value of this parameter to `true` if you cannot write to the keystore. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

-keyStoreStashFile

Set the value of this parameter to `true` if you want to create stash files for CMS type keystore. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

-enableCryptoOperations

Specifies if the keystore object will be used for hardware cryptographic operations or not. The default value is `false`. (Boolean, optional)

-keyStoreDescription

Specifies user defined text to describe the keystore of interest. (String, optional)

-keyStoreUsage

Specifies the keystore usage of interest. Specify `SSLKeys`, `KeySetKeys`, `RootKeys`, `DeletedKeys`, `DefaultSigners`, or `RSATokenKeys`. (String, optional)

-scopeName

The name that uniquely identifies the management scope, for example: `(cell):localhostNode01Cell`. (String, optional)

-controlRegionUser

Specifies the control region user to create a writable keystore object for the control regions key ring. Specify this option for SAF key rings when SAF writable key rings is enabled. (String, optional)

-servantRegionUser

Specifies the servant region user to create a writable keystore object for the servant regions key ring. Specify this option for SAF key rings when SAF writable key rings is enabled. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createKeyStore {-keyStoreName testKS -keyStoreType JCEKS
-keyStoreLocation c:\temp\testKeyFile.p12 -keyStorePassword testpwd
-keyStorePasswordVerify testpwd -keyStoreIsFileBased true -keyStoreInitAtStartup
true -keyStoreReadOnly false}
```

- Using Jython string:

```
AdminTask.createKeyStore(['-keyStoreName testKS -keyStoreType JCEKS -keyStoreLocation
c:\temp\testKeyFile.p12 -keyStorePassword testpwd -keyStorePasswordVerify testpwd
-keyStoreIsFileBased true -keyStoreInitAtStartup true -keyStoreReadOnly false'])
```

- Using Jython list:

```
AdminTask.createKeyStore(['-keyStoreName', 'testKS', '-keyStoreLocation', '-keyStoreType',
'JCEKS', 'c:\temp\testKeyFile.p12', '-keyStorePassword', 'testpwd',
'-keyStorePasswordVerify', 'testpwd', '-keyStoreIsFileBased', 'true',
'-keyStoreInitAtStartup', 'true', '-keyStoreReadOnly', 'false'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createKeyStore {-interactive}
```

- Using Jython:

```
AdminTask.createKeyStore('-interactive')
```

createCMSKeyStore

The **createCMSKeyStore** command creates a CMS keystore database and the keystore settings in the configuration.

Required parameters

-cmsKeyStoreURI

The URI of the CMS keystore. (String, required)

-pluginHostName

The host name of the plug-in. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createCMSKeyStore {-cmsKeyStoreURI CMSKeystoreURI -pluginHostName myHostName}
```

- Using Jython string:

```
AdminTask.createCMSKeyStore('-cmsKeyStoreURI CMSKeystoreURI -pluginHostName myHostName')
```

- Using Jython list:

```
AdminTask.createCMSKeyStore(['-cmsKeyStoreURI', 'CMSKeystoreURI', '-pluginHostName',  
'myHostName'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createCMSKeyStore {-interactive}
```

- Using Jython:

```
AdminTask.createCMSKeyStore('-interactive')
```

deleteKeyStore

The **deleteKeyStore** command deletes the settings of a keystore from the configuration and the keystore file.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore that you want to delete. (String, required)

Optional parameters

-scopeName

The name that uniquely identifies the management scope, for example: (cell):localhostNode01Cell.
(String, optional)

-removeKeyStoreFile

Specifies whether to remove the keystore file. Specify `true` to remove the keystore file or `false` to keep the keystore file in your configuration. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyStore {-keyStoreName testKS}
```

- Using Jython string:

```
AdminTask.deleteKeyStore(['-keyStoreName testKS'])
```

- Using Jython list:

```
AdminTask.deleteKeyStore(['-keyStoreName', 'testKS'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyStore {-interactive}
```

- Using Jython:

```
AdminTask.deleteKeyStore('-interactive')
```

exchangeSigners

The **exchangeSigners** command exchange signer certificate between keystores.

Required parameters

-keyStoreName1

The name that uniquely identifies a keystore. You must specify a second keystore name using the `keyStoreName2` parameter. (String, required)

-keyStoreScope1

The scope name of the keystore that you specified with the `keyStoreName1` parameter. (String, required)

-keyStoreName2

The name that uniquely identifies a keystore. You must specify a second keystore name using the `keyStoreName1` parameter. (String, required)

-keyStoreScope2

The scope name of the keystore that you specified with the `keyStoreName2` parameter. (String, required)

Optional parameters

-certificateAliasList1

A list of aliases separated by a comma. (String, optional)

-certificateAliasList2

A list of aliases separated by a comma. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask exchangeSigners {-keyStoreName1 testKS -certificateAliasList1 testCert1  
-keyStoreName2 secondKS -certificateAliasList2 certAlias}
```

- Using Jython string:

```
AdminTask.exchangeSigners('[-keyStoreName1 testKS -certificateAliasList1 testCert1  
-keyStoreName2 secondKS -certificateAliasList2 certAlias]')
```

- Using Jython list:

```
AdminTask.exchangeSigners(['-keyStoreName1', 'testKS', '-certificateAliasList1',  
'testCert1', '-keyStoreName2', 'secondKS', '-certificateAliasList2',  
'certAlias'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask exchangeSigners {-interactive}
```

- Using Jython:

```
AdminTask.exchangeSigners('-interactive')
```

getKeyStoreInfo

The **getKeyStoreInfo** command displays the settings of a particular keystore.

Required parameters

-name

The name that uniquely identifies the keystore. (String, required)

Optional parameters

-scopeName

The name that uniquely identifies the management scope, for example: (cell):localhostNode01Cell.
(String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getKeyStoreInfo {-name testKS}
```

- Using Jython string:

```
AdminTask.getKeyStoreInfo('[-name testKS]')
```

- Using Jython list:

```
AdminTask.getKeyStoreInfo(['-name', 'testKS'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getKeyStoreInfo {-interactive}
```

- Using Jython:

```
AdminTask.getKeyStoreInfo('-interactive')
```

listKeyFileAliases

The **listKeyFileAliases** command lists the certificates in a keystore file.

Required parameters

-keyFilePath

The path of the key file. (String, required)

-keyFilePassword

The password for the key file. (String, required)

-keyFileType

The key file type. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl: Windows

```
$AdminTask listKeyFileAliases {-keyFilePath c:\temp\testKeyFile.p12  
-keyFilePassword testPwd -keyFileType PKCS12}
```

HP-UX

Solaris

AIX

Linux

```
$AdminTask listKeyFileAliases {-keyFilePath /temp/testKeyFile.p12  
-keyFilePassword testPwd -keyFileType PKCS12}
```

- Using Jython string: Windows

```
AdminTask.listKeyFileAliases('[-keyFilePaht c:\temp\testKeyFile.p12  
-keyFilePassword testPwd -keyFileType PKCS12]')
```

HP-UX

Solaris

AIX

Linux

```
AdminTask.listKeyFileAliases('[-keyFilePaht /temp/testKeyFile.p12  
-keyFilePassword testPwd -keyFileType PKCS12]')
```

- Using Jython list: Windows

```
AdminTask.listKeyFileAliases(['-keyFilePaht', 'c:\temp\testKeyFile.p12',
'-keyFilePassword', 'testPwd', '-keyFileType', 'PKCS12'])
```

HP-UX
Solaris
AIX
Linux

```
AdminTask.listKeyFileAliases(['-keyFilePaht', '/temp/testKeyFile.p12',
'-keyFilePassword', 'testPwd', '-keyFileType', 'PKCS12'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listKeyFileAliases {-interactive}
```

- Using Jython:

```
AdminTask.listKeyFileAliases('-interactive')
```

listKeyStores

The **listKeyStores** command lists the keystore for a particular scope.

Required parameters

None.

Optional parameters

-displayObjectName

Specify the value of this parameter as `true` to list the keystore configuration objects within a scope. Set the value of this parameter to `false` to list the strings that contain the keystore name and management scope. (Boolean, optional)

-scopeName

Specifies the name that uniquely identifies the management scope, for example: `(cell):localhostNode01Cell`. (String, optional)

-all

Specify the value of this parameter as `true` to list all keystores. This parameter overrides the `scopeName` parameter. The default value is `false`. (Boolean, optional)

-keyStoreUsage

Specifies the keystore usage of interest. Specify `SSLKeys`, `KeySetKeys`, `RootKeys`, `DeletedKeys`, `DefaultSigners`, or `RSATokenKeys`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listKeyStores
```

- Using Jython:

```
AdminTask.listKeyStores()
```

Interactive mode example usage:

listKeyStoreTypes

The **listKeyStoreTypes** command lists all valid keystore types.

Required parameters

None.

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listKeyStoreTypes
```

- Using Jython:

```
AdminTask.listKeyStoreTypes()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listKeyStoreTypes {-interactive}
```

- Using Jython string:

```
AdminTask.listKeyStoreTypes('-interactive')
```

modifyKeyStore

The **modifyKeyStore** command modifies attributes for an existing keystore. Only some keystore attributes are modifiable, depending on what you are modifying. Use the following guidelines to use the command:

- To use this command to change the keystore file that the keystore object references, specify the `keyStoreName`, `keyStoreLocation`, `keyStoreType`, and `keyStorePassword` parameters.

-

Required parameters

-keyStoreName

Specifies the unique name that identifies the keystore. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the keystore. (String, optional)

-keyStoreType

Specifies one of the predefined keystore types. Valid values are JCEKS, CMSKS, PKCS12, PKCS11, and JKS. (String, optional)

-keyStoreLocation

Specifies the fully qualified location of the keystore file. To modify the location of the keystore file, you must specify the `keyStoreLocation`, `keyStoreType`, `keyStorePassword`, and `keyStoreName` parameters. (String, optional)

-keyStorePassword

Specifies the password to open the keystore. Use the `changeKeystorePassword` command to change the password of the keystore. (String, optional)

-keyStoresFileBased

Specifies whether the keystore is file based. To modify whether the keystore is file-based, specify the `keyStoresFileBased` and `keyStoreName` parameters. (Boolean, optional)

-keyStoreInitAtStartup

Specifies whether the keystore initiates at server startup. To modify whether the keystore initiates at server startup, specify the `keyStoreInitAtStartup` and `keyStoreName` parameters. (Boolean, optional)

-keyStoreReadOnly

Specifies whether the keystore is writable. To modify whether the keystore is read-only, specify the `keyStoreReadOnly` and `keyStoreName` parameters. (Boolean, optional)

-keyStoreDescription

Specifies a statement that describes the keystore. To modify the keystore description, specify the `keyStoreDescription` and `keyStoreName` parameters. (String, optional)

-keyStoreUsage

Specifies the keystore usage of interest. Specify `SSLKeys`, `KeySetKeys`, `RootKeys`, `DeletedKeys`, `DefaultSigners`, or `RSATokenKeys`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifyKeyStore {-keyStoreName CellDefaultKeyStore  
-keyStoreLocation c:\temp\testKeyFile.p12 -keyStoreType JCEKS  
-keyStorePassword my1password}
```

- Using Jython:

```
AdminTask.modifyKeyStore('-keyStoreName CellDefaultKeyStore -keyStoreLocation  
c:\temp\testKeyFile.p12 -keyStoreType JCEKS -keyStorePassword my1password')
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifyKeyStore {-interactive}
```

- Using Jython:

```
AdminTask.modifyKeyStore('-interactive')
```

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

“Creating an SSL configuration at the node scope using scripting” on page 957

An Secure Socket Layer (SSL) configuration references many other configuration objects. To help you make valid selections for the new SSL configuration before you create it, view information about existing configuration objects. Information about existing objects is also useful when you create a node scoped SSL configuration using the **createSSLConfig** command of the AdminTask object.

SSLConfigCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the SSLConfigCommands group can be used to create and manage Secure Sockets Layer (SSL) configurations and properties.

The SSLConfigCommands command group for the AdminTask object includes the following commands:

- “createSSLConfig”
- “createSSLConfigProperty” on page 992
- “deleteSSLConfig” on page 993
- “getSSLConfig” on page 994
- “getSSLConfigProperties” on page 995
- “listSSLCiphers” on page 996
- “listSSLConfigs” on page 996
- “listSSLConfigProperties” on page 997
- “listSSLRepertoires” on page 998
- “modifySSLConfig” on page 999

createSSLConfig

The createSSLConfig command creates an SSL configuration that is based on key store and trust store settings. You can use the SSL configuration settings to make the SSL connections.

Target object

None.

Required parameters

-alias

The name of the alias. (String, required)

-trustStoreNames

The key store that holds trust information used to validate the trust from remote connections. (String, required)

-keyStoreName

The key store that holds the personal certificates that provide identity for the connection. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

-clientKeyAlias

The certificate alias name for the client. (String, optional)

-serverKeyAlias

The certificate alias name for the server. (String, optional)

-type

The type of SSL configuration. (String, optional)

-clientAuthentication

Set the value of this parameter to true to request client authentication. Otherwise, set the value of this parameter to false. (Boolean, optional)

-securityLevel

The cipher group that you want to use. Valid values include: HIGH, MEDIUM, LOW, and CUSTOM. (String, optional)

-enabledCiphers

A list of ciphers used during SSL handshake. (String, optional)

-jsseProvider

One of the JSSE providers. (String, optional)

-clientAuthenticationSupported

Set the value of this parameter to true to support client authentication. Otherwise, set the value of this parameter to false. (Boolean, optional)

-sslProtocol

The protocol type for the SSL handshake. Valid values include: SSL_TLS, SSL, SSLv2, SSLv3, TLS, TLSv1. (String, optional)

-trustManagerObjectName

A list of trust managers separated by commas. (String, optional)

-trustStoreScopeName

The management scope name of the trust store. (String, optional)

-keyStoreScopeName

The management scope name of the key store. (String, optional)

-ssslKeyRingName

Specifies a system SSL (SSSL) key ring name. The value for this parameter has no affect unless the SSL configuration type is SSSL. (String, optional)

Example output

The command returns the configuration object name of the new SSL configuration object.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createSSLConfig {-alias testSSLCfg -clientKeyAlias key1
-serverKeyAlias key2 -trustStoreNames trustKS -keyStoreName
testKS -keyManagerName testKeyMgr}
```

- Using Jython string:

```
AdminTask.createSSLConfig('[-alias testSSLCfg -clientKeyAlias key1
-serverKeyAlias key2 -trustStoreNames trustKS -keyStoreName
testKS -keyManagerName testKeyMgr]')
```

- Using Jython list:

```
AdminTask.createSSLConfig(['-alias', 'testSSLCfg', '-clientKeyAlias',
'key1', '-serverKeyAlias', 'key2', '-trustStoreNames', 'trustKS',
'-keyStoreName', 'testKS', '-keyManagerName', 'testKeyMgr'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createSSLConfig {-interactive}
```

- Using Jython:

```
AdminTask.createSSLConfig('-interactive')
```

createSSLConfigProperty

The createSSLConfigProperty command creates a property for an SSL configuration. Use this command to set SSL configuration settings that are different than the settings in the SSL configuration object.

Target object

None.

Required parameters

-sslConfigAliasName

The alias name of the SSL configuration. (String, required)

-propertyName

The name of the property. (String, required)

-propertyValue

The value of the property. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createSSLConfigProperty {-sslConfigAliasName NodeDefaultSSLSettings  
-scopeName (cell):localhostNode01Cell:(node):localhostNode01 -propertyName  
test.property -propertyValue testValue}
```

- Using Jython string:

```
AdminTask.createSSLConfigProperty(['-sslConfigAliasName NodeDefaultSSLSettings  
-scopeName (cell):localhostNode01Cell:(node):localhostNode01 -propertyName  
test.property -propertyValue testValue'])
```

- Using Jython list:

```
AdminTask.createSSLConfigProperty(['-sslConfigAliasName', 'NodeDefaultSSLSettings',  
'-scopeName', '(cell):localhostNode01Cell:(node):localhostNode01', '-propertyName',  
'test.property', '-propertyValue', 'testValue'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createSSLConfigProperty {-interactive}
```

- Using Jython:

```
AdminTask.createSSLConfigProperty('-interactive')
```

deleteSSLConfig

The deleteSSLConfig command deletes the SSL configuration object that you specify from the configuration.

Target object

None.

Required parameters and return values

-alias

The name of the alias. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteSSLConfig {-alias NodeDefaultSSLSettings -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01}
```

- Using Jython string:

```
AdminTask.deleteSSLConfig(['-alias NodeDefaultSSLSettings -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01'])
```

- Using Jython list:

```
AdminTask.deleteSSLConfig(['-alias', 'NodeDefaultSSLSettings', '-scopeName',  
'(cell):localhostNode01Cell:(node):localhostNode01'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteSSLConfig {-interactive}
```

- Using Jython:

```
AdminTask.deleteSSLConfig('-interactive')
```

getSSLConfig

The getSSLConfig command obtains information about an SSL configuration and displays the settings.

Target object

None.

Required parameters and return values

-alias

The name of the alias. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

Example output

The command returns information about the SSL configuration of interest.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfig {-alias NodeDefaultSSLSettings -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01}
```

- Using Jython string:

```
AdminTask.getSSLConfig(['-alias NodeDefaultSSLSettings -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01'])
```

- Using Jython list:

```
AdminTask.getSSLConfig(['-alias', 'NodeDefaultSSLSettings', '-scopeName',  
'(cell):localhostNode01Cell:(node):localhostNode01'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfig {-interactive}
```

- Using Jython:

```
AdminTask.getSSLConfig('-interactive')
```

getSSLConfigProperties

The `getSSLConfigProperties` command obtains information about SSL configuration properties.

Target object

None.

Required parameters and return values

-alias

The name of the alias. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

Example output

The command returns additional information about the SSL configuration properties.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfigProperties {-sslConfigAliasName NodeDefaultSSLSettings  
-scopeName (cell):localhostNode01Cell:(node):localhostNode01}
```

- Using Jython string:

```
AdminTask.getSSLConfigProperties(['-sslConfigAliasName NodeDefaultSSLSettings  
-scopeName (cell):localhostNode01Cell:(node):localhostNode01'])
```

- Using Jython list:

```
AdminTask.getSSLConfigProperties(['-sslConfigAliasName', 'NodeDefaultSSLSettings',  
'-scopeName', '(cell):localhostNode01Cell:(node):localhostNode01'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfigProperties {-interactive}
```

- Using Jython:

```
AdminTask.getSSLConfigProperties('-interactive')
```

listSSLCiphers

The listSSLCiphers command lists the SSL ciphers.

Target object

None.

Required parameters

-sslConfigAliasName

The alias name of the SSL configuration. (String, required)

-securityLevel

The cipher group that you want to use. Valid values include: HIGH, MEDIUM, LOW, and CUSTOM. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

-provider

(String, optional)

Example output

The command returns a list of SSL ciphers.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listSSLCiphers {-sslConfigAliasName testSSLCfg  
-securityLevel HIGH}
```
- Using Jython string:

```
AdminTask.listSSLCiphers(['-sslConfigAliasName testSSLCfg  
-securityLevel HIGH'])
```
- Using Jython list:

```
AdminTask.listSSLCiphers(['-sslConfigAliasName', 'testSSLCfg',  
'-securityLevel', 'HIGH'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listSSLCiphers {-interactive}
```
- Using Jython:

```
AdminTask.listSSLCiphers('-interactive')
```

listSSLConfigs

The listSSLConfigs command lists the defined SSL configurations within a management scope.

Target object

None.

Optional parameters

-scopeName

The name of the scope. (String, optional)

-displayObjectName

Set the value of this parameter to `true` to list the SSL configuration objects within the scope. Set the value of this parameter to `false` to list the strings that contain the SSL configuration alias and management scope. (Boolean, optional)

-all

Specify the value of this parameter as `true` to list all SSL configurations. This parameter overrides the `scopeName` parameter. The default value is `false`. (Boolean, optional)

Example output

The command returns a list of defined SSL configurations.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigs {-scopeName (cell):localhostNode01Cell:(node):localhostNode01  
-displayObjectName true}
```

- Using Jython string:

```
AdminTask.listSSLConfigs(['-scopeName (cell):localhostNode01Cell:(node):localhostNode01  
-displayObjectName true'])
```

- Using Jython list:

```
AdminTask.listSSLConfigs(['-scopeName', '(cell):localhostNode01Cell:(node):localhostNode01',  
'-displayObjectName', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigs {-interactive}
```

- Using Jython:

```
AdminTask.listSSLConfigs('-interactive')
```

listSSLConfigProperties

The `listSSLConfigProperties` command lists the properties for a SSL configuration.

Target object

None.

Required parameters

-alias

The alias name of the SSL configuration. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

-displayObjectName

Set the value of this parameter to `true` to list the SSL configuration objects within the scope. Set the value of this parameter to `false` to list the strings that contain the SSL configuration alias and management scope. (Boolean, optional)

Example output

The command returns SSL configuration properties.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigProperty {-alias SSL123 -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01 -displayObjectName true}
```
- Using Jython string:

```
AdminTask.listSSLConfigProperty(['-alias SSL123 -scopeName  
(cell):localhostNode01Cell:(node):localhostNode01 -displayObjectName true'])
```
- Using Jython list:

```
AdminTask.listSSLConfigProperty(['-alias', 'SSL123', '-scopeName',  
'(cell):localhostNode01Cell:(node):localhostNode01', '-displayObjectName', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigProperties {-interactive}
```
- Using Jython:

```
AdminTask.listSSLConfigProperties('-interactive')
```

listSSLRepertoires

The `listSSLRepertoires` command lists all of the Secure Sockets Layer (SSL) configuration instances that you can associate with an SSL inbound channel. If you create a new SSL alias using the administrative console, the alias name is automatically created in the `node_name/alias_name` format. However, if you create a new SSL alias using the `wsadmin` tool, you must create the SSL alias and specify both the node name and alias name in the `node_name/alias_name` format.

Target object

SSLInboundChannel instance for which the SSLConfig candidates are listed.

Required parameters

None.

Optional parameters

None.

Sample output

The command returns a list of eligible SSL configuration object names.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listSSLRepertoires SSL_3(cells/mybuildCell01/nodes/mybuildNode01/servers/
server2|server.xml#SSLInboundChannel_1093445762330)
```

- Using Jython string:

```
print AdminTask.listSSLRepertoires('SSL_3(cells/mybuildCell01/nodes/mybuildNode01/
servers/server2|server.xml#SSLInboundChannel_1093445762330)')
```

- Using Jython list:

```
print AdminTask.listSSLRepertoires('SSL_3(cells/mybuildCell01/nodes/mybuildNode01/
servers/server2|server.xml#SSLInboundChannel_1093445762330)')
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listSSLRepertoires {-interactive}
```

- Using Jython:

```
print AdminTask.listSSLRepertoires('-interactive')
```

modifySSLConfig

The modifySSLConfig command modifies the settings of an existing SSL configuration.

Target object

None.

Required parameters

-alias

The name of the alias. (String, required)

Optional parameters

-scopeName

The name of the scope. (String, optional)

-clientKeyAlias

The certificate alias name for the client. (String, optional)

-serverKeyAlias

The certificate alias name for the server. (String, optional)

-type

The type of SSL configuration. (String, optional)

-clientAuthentication

Set the value of this parameter to true to request client authentication. Otherwise, set the value of this parameter to false. (Boolean, optional)

-securityLevel

The cipher group that you want to use. Valid values include: HIGH, MEDIUM, LOW, and CUSTOM. (String, optional)

-enabledCiphers

A list of ciphers used during SSL handshake. (String, optional)

-jsseProvider

One of the JSSE providers. (String, optional)

-clientAuthenticationSupported

Set the value of this parameter to true to support client authentication. Otherwise, set the value of this parameter to false. (Boolean, optional)

-sslProtocol

The protocol type for the SSL handshake. Valid values include: SSL_TLS, SSL, SSLv2, SSLv3, TLS, TLSv1. (String, optional)

-trustManagerObjectNames

A list of trust managers separated by commas. (String, optional)

-trustStoreNames

The key store that holds trust information used to validate the trust from remote connections. (String, optional)

-trustStoreScopeName

The management scope name of the trust store. (String, optional)

-keyStoreName

The key store that holds the personal certificates that provide identity for the connection. (String, optional)

-keyStoreScopeName

The management scope name of the key store. (String, optional)

-ssslKeyRingName

Specifies a system SSL (SSSL) key ring name. The value for this parameter has no affect unless the SSL configuration type is SSSL. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifySSLConfig {-alias testSSLCfg -clientKeyAlias tstKey1
-serverKeyAlias tstKey2 -securityLevel LOW}
```

- Using Jython string:

```
AdminTask.modifySSLConfig(['-alias testSSLCfg -clientKeyAlias tstKey1
-serverKeyAlias tstKey2 -securityLevel LOW'])
```

- Using Jython list:

```
AdminTask.modifySSLConfig(['-alias', 'testSSLCfg', '-clientKeyAlias', 'tstKey1',
'-serverKeyAlias', 'tstKey2', '-securityLevel', 'LOW'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifySSLConfig {-interactive}
```

- Using Jython:

```
AdminTask.modifySSLConfig('-interactive')
```

SSLConfigGroupCommands group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the SSLConfigGroupCommands group can be used to create and manage SSL configuration groups.

The SSLConfigGroupCommands command group for the AdminTask object includes the following commands:

- “deleteSSLConfigGroup” on page 1001

- “getSSLConfigGroup”
- “listSSLConfigGroups” on page 1002
- “modifySSLConfigGroup” on page 1003

deleteSSLConfigGroup

The deleteSSLConfigGroup command deletes a SSL configuration group from the configuration.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the SSL configuration group. (String, required)

-direction

Specifies the direction to which the SSL configuration applies. Valid values include inbound or outbound. (String, required)

Optional parameters

-scopeName

Specifies the name that uniquely identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteSSLConfigGroup {-name createSSLCfgGrp -direction inbound}
```

- Using Jython string:

```
AdminTask.deleteSSLConfigGroup(['-name createSSLCfgGrp -direction inbound'])
```

- Using Jython list:

```
AdminTask.deleteSSLConfigGroup(['-name', 'createSSLCfgGrp', '-direction', 'inbound'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteSSLConfigGroup {-interactive}
```

- Using Jython:

```
AdminTask.deleteSSLConfigGroup('-interactive')
```

getSSLConfigGroup

The getSSLConfigGroup command returns information about a SSL configuration setting.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the SSL configuration group. (String, required)

-direction

Specifies the direction to which the SSL configuration applies. Valid values include inbound or outbound. (String, required)

Optional parameters

-scopeName

Specifies the name that uniquely identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfigGroup {-name createSSLCfgGrp -direction inbound}
```

- Using Jython string:

```
AdminTask.getSSLConfigGroup('[-name createSSLCfgGrp -direction inbound]')
```

- Using Jython list:

```
AdminTask.getSSLConfigGroup(['-name', 'createSSLCfgGrp', '-direction', 'inbound'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getSSLConfigGroup {-interactive}
```

- Using Jython:

```
AdminTask.getSSLConfigGroup('-interactive')
```

listSSLConfigGroups

The listSSLConfigGroups command lists the SSL configuration groups within a scope and a direction.

Target object

None.

Required parameters

None.

Optional parameters

-direction

Specifies the direction to which the SSL configuration applies. Valid values include inbound or outbound. (String, optional)

-scopeName

Specifies the name that uniquely identifies the management scope. (String, optional)

-displayObjectName

If you set this parameter to true, the command returns a list of all of the SSL configuration group objects within the scope. If you set this parameter to false, the command returns a list of strings that contain the SSL configuration name and management scope. (Boolean, optional)

-all

Specify the value of this parameter as true to list all SSL configuration groups. This parameter overrides the scopeName parameter. The default value is false. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigGroups {-displayObjectName true}
```

- Using Jython string:

```
AdminTask.listSSLConfigGroups(['-displayObjectName true'])
```

- Using Jython list:

```
AdminTask.listSSLConfigGroups(['-displayObjectName' 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listSSLConfigGroups {-interactive}
```

- Using Jython:

```
AdminTask.listSSLConfigGroups('-interactive')
```

modifySSLConfigGroup

The `modifySSLConfigGroup` command modifies the setting of an existing SSL configuration group.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the SSL configuration group. (String, required)

-direction

Specifies the direction to which the SSL configuration applies. Valid values include `inbound` or `outbound`. (String, required)

Optional parameters

-certificateAlias

Specifies a unique name to identify a certificate. (String, optional)

-scopeName

Specifies the name that uniquely identifies the management scope. (String, optional)

-sslConfigAliasName

Specifies the alias that uniquely identifies the SSL configurations in the group. (String, optional)

-sslConfigScopeName

Specifies the scope that uniquely identifies the SSL configurations in the group. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifySSLConfigGroup {-name createSSLCfgGrp -direction inbound -certificateAlias alias2}
```

- Using Jython string:

```
AdminTask.modifySSLConfigGroup(['-name createSSLCfgGrp -direction inbound -certificateAlias alias2'])
```

- Using Jython list:

```
AdminTask.modifySSLConfigGroup(['-name', 'createSSLCfgGrp', '-direction', 'inbound', '-certificateAlias', 'alias2'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifySSLConfigGroup {-interactive}
```

- Using Jython:

```
AdminTask.modifySSLConfigGroup('-interactive')
```

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

“Creating an SSL configuration at the node scope using scripting” on page 957

An Secure Socket Layer (SSL) configuration references many other configuration objects. To help you make valid selections for the new SSL configuration before you create it, view information about existing configuration objects. Information about existing objects is also useful when you create a node scoped SSL configuration using the **createSSLConfig** command of the AdminTask object.

TrustManagerCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the TrustManagerCommands group can be used to create, delete, and query trust manager settings in your configuration. You can also use these commands to create a custom trust manager for a pure client.

The TrustManagerCommands command group for the AdminTask object includes the following commands:

- “deleteTrustManager”
- “getTrustManager” on page 1005
- “listTrustManagers” on page 1006
- “modifyTrustManager” on page 1006

deleteTrustManager

The **deleteTrustManager** command deletes the trust manager settings from the configuration.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the trust manager. (String, required)

Optional parameters

-scopeName

Specifies the name of the scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteTrustManager {-name testTM}`
- Using Jython string:
`AdminTask.deleteTrustManager(['-name testTM'])`
- Using Jython list:
`AdminTask.deleteTrustManager(['-name', 'testTM'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteTrustManager {-interactive}`
- Using Jython:
`AdminTask.deleteTrustManager('-interactive')`

getTrustManager

The **getTrustManager** command obtains the setting of a trust manager.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the trust manager. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask getTrustManager {-name testTM}`
- Using Jython string:
`AdminTask.getTrustManager(['-name testTM'])`
- Using Jython list:
`AdminTask.getTrustManager(['-name', 'testTM'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask getTrustManager {-interactive}`
- Using Jython:
`AdminTask.getTrustManager('-interactive')`

listTrustManagers

The **listTrustManagers** command lists the trust managers within a particular management scope.

Target object

None.

Required parameters

None.

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-displayObjectName

Set the value of this parameter to true to list the trust manager objects within a scope. Set the value of this parameter to false to list the strings that contain the trust manager name and management scope. (Boolean, optional)

-all

Specify the value of this parameter as true to list all trust managers. This parameter overrides the scopeName parameter. The default value is false. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listTrustManagers {-displayObjectName true}
```
- Using Jython string:

```
AdminTask.listTrustManagers(['-displayObjectName true'])
```
- Using Jython list:

```
AdminTask.listTrustManagers(['-displayObjectName', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listTrustManagers {-interactive}
```
- Using Jython:

```
AdminTask.listTrustManagers('-interactive')
```

modifyTrustManager

The **modifyTrustManager** command changes existing trust manager settings.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the trust manager. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-provider

Specifies the provider name of the trust manager. (String, optional)

-algorithm

Specifies the algorithm name of the trust manager. (String, optional)

-trustManagerClass

Specifies a class that implements the `javax.net.ssl.X509TrustManager` interface. You cannot use this parameter with the provider or algorithm parameters. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifyTrustManager {-name testTM -trustManagerClass test.trust.manager}
```

- Using Jython string:

```
AdminTask.modifyTrustManager(['-name testTM -trustManagerClass test.trust.manager'])
```

- Using Jython list:

```
AdminTask.modifyTrustManager(['-name', 'testTM', '-trustManagerClass', 'test.trust.manager'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifyTrustManager {-interactive}
```

- Using Jython:

```
AdminTask.modifyTrustManager('-interactive')
```

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

“Creating a custom trust manager configuration” on page 646

You can create a custom trust manager configuration at any management scope and associate the new trust manager with a Secure Sockets Layer (SSL) configuration.

KeySetCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the KeySetCommands group can be used to create, delete, and query for key set settings in your configuration.

The `KeySetCommands` command group for the `AdminTask` object includes the following commands:

- “`createKeySet`”
- “`deleteKeySet`” on page 1009
- “`generateKeyForKeySet`” on page 1010
- “`getKeySet`” on page 1011
- “`listKeySets`” on page 1011
- “`modifyKeySet`” on page 1012

createKeySet

The **`createKeySet`** command creates the key set settings in the configuration. Use this command to control key instances that have the same type.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set. (String, required)

-aliasPrefix

Specifies the prefix for the key alias when a new key generates. (String, required)

-password

Specifies the password that protects the key in the keystore. (String, required)

-maxKeyReferences

Specifies the maximum number of key references from the returned keys in the key set of interest. (Integer, required)

-keyStoreName

Specifies the keystore that contains the keys. (String, required)

Optional parameters

-scopeName

Specifies the unique name of the management scope. (String, optional)

-deleteOldKeys

Set the value of this parameter to `true` to delete old keys when new keys are generated. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

-keyGenerationClass

Specifies the class that is used to generate new keys in the key set. (String, optional)

-keyStoreScopeName

Specifies the management scope where the keystore is located. (String, optional)

-isKeyPair

Set the value of this parameter to `true` if the keys in the key set are key pairs. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

Example output

The command returns the configuration object name of the key set object that you created.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createKeySet {-name testKeySet -aliasPrefix test -password pwd  
-maxKeyReferences 2 -deleteOldKeys true -keyStoreName testKeyStore -isKeyPair false}
```

- Using Jython string:

```
AdminTask.createKeySet(['-name testKeySet -aliasPrefix test -password pwd  
-maxKeyReferences 2 -deleteOldKeys true -keyStoreName testKeyStore -isKeyPair false'])
```

- Using Jython list:

```
AdminTask.createKeySet(['-name', 'testKeySet', '-aliasPrefix', 'test',  
'-password', 'pwd', '-maxKeyReferences', '2', '-deleteOldKeys', 'true',  
'-keyStoreName', 'testKeyStore', '-isKeyPair', 'false'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createKeySet {-interactive}
```

- Using Jython string:

```
AdminTask.createKeySet ('[-interactive]')
```

- Using Jython list:

```
AdminTask.createKeySet (['-interactive'])
```

deleteKeySet

The **deleteKeySet** command deletes the settings of a key set from the configuration.

Target object

None.

Required parameters

-name

The name that uniquely identifies the key set. (String, required)

Optional parameters

-scopeName

Specifies the unique name of the management scope. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteKeySet{ -name testKeySet}
```

- Using Jython string:

```
AdminTask.deleteKeySet(['-name testKeySet'])
```

- Using Jython list:

```
AdminTask.deleteKeySet(['-name', 'testKeySet'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteKeySet {-interactive}
```

- Using Jython string:
`AdminTask.deleteKeySet ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteKeySet (['-interactive'])`

generateKeyForKeySet

The **generateKeyForKeySet** command generates keys for the keys in the key set.

Target object

None.

Required parameters

- keySetName**
Specifies the name of the key set. (String, required)

Optional parameters

- keySetScope**
Specifies the scope of the key set. (String, optional)
- keySetSaveConfig**
Set the value of this parameter to true to save the configuration of the key set. Otherwise, set the value of this parameter to false. (Boolean, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask generateKeyForKeySet{ -keySetName testKeySet }`
- Using Jython string:
`AdminTask.generateKeyForKeySet(['-keySetName testKeySet'])`
- Using Jython list:
`AdminTask.generateKeyForKeySet(['-keySetName', 'testKeySet'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask generateKeyForKeySet {-interactive}`
- Using Jython string:
`AdminTask.generateKeyForKeySet ('[-interactive]')`
- Using Jython list:
`AdminTask.generateKeyForKeySet (['-interactive'])`

getKeySet

The **getKeySet** command displays the settings of a particular key set.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set. (String, required)

Optional parameters

-scopeName

Specifies the unique name of the management scope. (String, optional)

Example output

The command returns the settings of the specified key set group.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask getKeySet {-name testKeySet}`
- Using Jython string:
`AdminTask.getKeySet ('[-name testKeySet]')`
- Using Jython list:
`AdminTask.getKeySet (['-name', 'testKeySet'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask getKeySet {-interactive}`
- Using Jython string:
`AdminTask.getKeySet (['-interactive'])`
- Using Jython list:
`AdminTask.getKeySet (['-interactive'])`

listKeySets

The **listKeySets** command lists the key sets in a particular scope.

Target object

None.

Required parameters

None.

Optional parameters

-scopeName

Specifies the unique name of the management scope. (String, optional)

-displayObjectNames

Set the value of this parameter to `true` to list the key set configuration objects within the scope. Set the value of this parameter to `false` if you want to list the strings that contain the key set group name and management scope. (Boolean, optional)

-all

Specify the value of this parameter as `true` to list all key sets. This parameter overrides the `scopeName` parameter. The default value is `false`. (Boolean, optional)

Example output

The command returns the key sets for the scope that you specified.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listKeySets {-displayObjectName true}
```
- Using Jython string:

```
AdminTask.listKeySets ('[-displayObjectName true]')
```
- Using Jython list:

```
AdminTask.listKeySets (['-displayObjectName', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listKeySets {-interactive}
```
- Using Jython string:

```
AdminTask.listKeySets ('[-interactive]')
```
- Using Jython list:

```
AdminTask.listKeySets (['-interactive'])
```

modifyKeySet

The **modifyKeySet** command changes the settings of an existing key set.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set. (String, required)

Optional parameters

-scopeName

Specifies the unique name of the management scope. (String, optional)

-aliasPrefix

Specifies the prefix for the key alias when a new key generates. (String, optional)

- password**
Specifies the password that protects the key in the keystore. (String, optional)
- maxKeyReferences**
Specifies the maximum number of key references from the returned keys in the key set of interest. (Integer, optional)
- deleteOldKeys**
Set the value of this parameter to `true` to delete old keys when new keys are generated. Otherwise, set the value of this parameter to `false`. (Boolean, optional)
- keyGenerationClass**
Specifies the class that is used to generate new keys in the key set. (String, optional)
- keyStoreName**
Specifies the keystore that contains the keys. (String, optional)
- keyStoreScopeName**
Specifies the management scope where the keystore is located. (String, optional)
- isKeyPair**
Set the value of this parameter to `true` if the keys in the key set are key pairs. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask modifyKeySet {-name testKeySet -maxKeyReferences 3 -deleteOldKeys false}`
- Using Jython string:
`AdminTask.modifyKeySet ('[-name testKeySet -maxKeyReferences 3 -deleteOldKeys false]')`
- Using Jython list:
`AdminTask.modifyKeySet (['-name', 'testKeySet', '-maxKeyReferences', '3', '-deleteOldKeys', 'false'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask modifyKeySet {-interactive}`
- Using Jython string:
`AdminTask.modifyKeySet (['-interactive'])`
- Using Jython list:
`AdminTask.modifyKeySet (['-interactive'])`

KeyReferenceCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the KeyReferenceCommands group can be used to create and manage the key reference settings for key set objects in your configuration.

The KeyReferenceCommands command group for the AdminTask object includes the following commands:

- “createKeyReference” on page 1014

- “deleteKeyReference” on page 1015
- “getKeyReference” on page 1015
- “listKeyReferences” on page 1016

createKeyReference

The **createKeyReference** command creates the key reference setting in the configuration for key set objects.

Target object

None.

Required parameters and return values

-keySetName

The name that uniquely identifies the key set to which the key reference belongs. (String, required)

-keySetScope

The management scope of the key set. (String, optional)

-keyAlias

The alias name that identifies the key for the key set that you specify. (String, required)

-keyPassword

The password used for encrypting the key. (String, optional)

-keyPasswordVerify

The password used for encrypting the key. (String, optional)

-version

The version of the key reference. (String, optional)

-keyReferenceSaveConfig

Set the value of this parameter to `true` to save the key reference to the configuration. Otherwise, set the value to `false`. (String, optional)

- Returns: The configuration object name of the key reference scope object that you created.

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask createKeyReference {-keySetName testKeySet -keyAlias testKey
-pasword testPWD -passwordVerify testPWD -keyReferenceSaveConfig true}
```
- Using Jython string:


```
AdminTask.createKeyReference ('[-keySetName testKeySet -keyAlias testKey
-pasword testPWD -passwordVerify testPWD -keyReferenceSaveConfig true]')
```
- Using Jython list:


```
AdminTask.createKeyReference (['-keySetName', 'testKeySet', '-keyAlias', 'testKey',
'-pasword', 'testPWD', '-passwordVerify', 'testPWD', '-keyReferenceSaveConfig', 'true'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask createKeyReference {-interactive}
```
- Using Jython string:


```
AdminTask.createKeyReference ('[-interactive]')
```
- Using Jython list:

```
AdminTask.createKeyReference (['-interactive'])
```

deleteKeyReference

The **deleteKeyReference** command deletes a key reference object from the key set object in the configuration.

Target object

None.

Required parameters and return values

-keySetName

The name that uniquely identifies the key set to which the key reference belongs. (String, required)

-keySetScope

The management scope of the key set. (String, optional)

-keyAlias

The alias name that identifies the key for the key set that you specify. (String, required)

- Returns: None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyReference { -keySetName testKeySet -keyAlias testKey }
```
- Using Jython string:

```
AdminTask.deleteKeyReference (['-keySetName testKeySet -keyAlias testKey'])
```
- Using Jython list:

```
AdminTask.deleteKeyReference (['-keySetName', 'testKeySet', '-keyAlias', 'testKey'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteKeyReference {-interactive}
```
- Using Jython string:

```
AdminTask.deleteKeyReference (['-interactive'])
```
- Using Jython list:

```
AdminTask.deleteKeyReference (['-interactive'])
```

getKeyReference

The **getKeyReference** command displays the setting of a key reference object.

Target object

None.

Required parameters and return values

-keySetName

The name that uniquely identifies the key set to which the key reference belongs. (String, required)

-keySetScope

The management scope of the key set. (String, optional)

-keyAlias

The alias name that identifies the key for the key set that you specify. (String, required)

- Returns: The settings of the key reference object.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getKeyReference { -keySetName testKeySet -keyAlias testKey }
```
- Using Jython string:

```
AdminTask.getKeyReference ('[-keySetName testKeySet -keyAlias testKey]')
```
- Using Jython list:

```
AdminTask.getKeyReference (['-keySetName', 'testKeySet', '-keyAlias', 'testKey'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getKeyReference {-interactive}
```
- Using Jython string:

```
AdminTask.getKeyReference ('[-interactive]')
```
- Using Jython list:

```
AdminTask.getKeyReference (['-interactive'])
```

listKeyReferences

The **listKeyReferences** command lists the key references for a particular key set in the configuration.

Target object

None.

Required parameters and return values

-keySetName

The name that uniquely identifies the key set to which the key reference belongs. (String, required)

-keySetScope

The management scope of the key set. (String, optional)

- Returns: The configuration object name of the key reference scope object that you created.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listKeyReferences { -keySetName testKeySet }
```
- Using Jython string:

```
AdminTask.listKeyReferences ('[-keySetName testKeySet]')
```
- Using Jython list:

```
AdminTask.listKeyReferences (['-keySetName', 'testKeySet'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listKeyReferences {-interactive}
```

- Using Jython string:
AdminTask.listKeyReferences ('[-interactive]')
- Using Jython list:
AdminTask.listKeyReferences (['-interactive'])

KeySetGroupCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the KeySetGroupCommands group can be used to create and manage key set groups. Use these commands to manage groups of public, private, and shared keys.

The KeySetGroupCommands command group for the AdminTask object includes the following commands:

- “deleteKeySetGroup”
- “generateKeysForKeySetGroup”
- “getKeySetGroup” on page 1018
- “listKeySetGroups” on page 1018
- “modifyKeySetGroup” on page 1019

deleteKeySetGroup

The **deleteKeySetGroup** command deletes the settings of a key set group from the configuration.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set group. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

generateKeysForKeySetGroup

The **generateKeysForKeySetGroup** command generates keys for all of the keys in the key sets that make up the key set group.

Target object

None.

Required parameters

-keySetName

Specifies the name of the key set group. (String, required)

Optional parameters

-keySetGroupScope

Specifies the scope of the key set group. (String, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

getKeySetGroup

The **getKeySetGroup** command displays the settings of a particular key set group.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set group. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

listKeySetGroups

The **listKeySetGroups** command lists the key set groups for a particular scope.

Target object

None.

Required parameters

None.

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-displayObjectNames

If you set the value of this parameter to true, the command returns a list of all of the key set group objects within a scope. If you set the value of this parameter to false, the command returns a list of strings that contain the key set group name and management scope. (Boolean, optional)

-all

Specify the value of this parameter as `true` to list all key set groups. This parameter overrides the `scopeName` parameter. The default value is `false`. (Boolean, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

modifyKeySetGroup

The **modifyKeySetGroup** command changes the settings of an existing key set group.

Target object

None.

Required parameters

-name

Specifies the name that uniquely identifies the key set group. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-autoGenerate

Set the value of this parameter to `true` if you want to automatically generate keys. If not, set the value to `false`. (Boolean, optional)

-wsScheduleName

Specifies the name of the scheduler to use to perform key generation. (String, optional)

-keySetObjectNames

A list of key set configuration names separated by colons (:). (String, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Creating a key set group configuration” on page 722

A key set group manages one or more key sets. WebSphere Application Server uses key set groups to automatically generate cryptographic keys or multiple synchronized key sets.

Related reference

“Key set groups settings” on page 728

Use this page to create new key set groups.

DynamicSSLConfigSelections command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the DynamicSSLConfigSelections group can be used to create, delete, and query dynamic SSL configuration selection objects.

The DynamicSSLConfigSelections command group for the AdminTask object includes the following commands:

- “deleteDynamicSSLConfigSelection”
- “getDynamicSSLConfigSelection” on page 1021
- “listDynamicSSLConfigSelections” on page 1021

deleteDynamicSSLConfigSelection

The **deleteDynamicSSLConfigSelection** command deletes the dynamic SSL configuration selection from the configuration.

Target object

None.

Required parameters

-dynSSLConfigSelectionName

Specifies the name that uniquely identifies the dynamic SSL configuration selection. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteDynamicSSLConfigSelection {-dynSSLConfigSelectionName sampleConfigSelection}
```

- Using Jython:

```
AdminTask.deleteDynamicSSLConfigSelection(-dynSSLConfigSelectionName sampleConfigSelection)
```

•

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteDynamicSSLConfigSelection {-interactive}
```

- Using Jython:

```
AdminTask.deleteDynamicSSLConfigSelection('-interactive')
```

getDynamicSSLConfigSelection

The **getDynamicSSLConfigSelection** command obtains information about a particular dynamic SSL configuration selection.

Target object

None.

Required parameters

-dynSSLConfigSelectionName

Specifies the name that uniquely identifies the dynamic SSL configuration selection. (String, required)

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getDynamicSSLConfigSelection {-dynSSLConfigSelectionName sampleConfigSelection}
```

- Using Jython:

```
AdminTask.getDynamicSSLConfigSelection(-dynSSLConfigSelectionName sampleConfigSelection)
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getDynamicSSLConfigSelection {-interactive}
```

- Using Jython:

```
AdminTask.getDynamicSSLConfigSelection('-interactive')
```

listDynamicSSLConfigSelections

The **listDynamicSSLConfigSelections** command lists the configuration objects name for a dynamic SSL configuration selection.

Target object

None.

Required parameters

None.

Optional parameters

-scopeName

Specifies the unique name that identifies the management scope. (String, optional)

-all

Specify the value of this parameter as true to list all dynamic SSL configuration selections. This parameter overrides the scopeName parameter. The default value is false. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listDynamicSSLConfigSelections`
- Using Jython:
`AdminTask.listDynamicSSLConfigSelections()`
- Using Jacl:
`$AdminTask listDynamic SSLConfigSelections`
- Using Jython:
`AdminTask.listDynamic SSLConfigSelections()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listDynamicSSLConfigSelections {-interactive}`
- Using Jython:
`AdminTask.listDynamicSSLConfigSelections('-interactive')`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Replacing an existing self-signed certificate” on page 687

Occasionally, you need to replace an existing or expired self-signed certificate with a new certificate. Certificates are referenced in the runtime configuration by the Secure Sockets Layer (SSL) Configuration object and the Dynamic SSL Configuration Selection object. You can replace a certificate with a new certificate alias reference or with a new signer certificate.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

PersonalCertificateCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the PersonalCertificateCommands group can be used to create and manage personal or signer certificates.

The PersonalCertificateCommands command group for the AdminTask object includes the following commands:

- “createChainedCertificate”
- “createSelfSignedCertificate” on page 1024
- “deleteCertificate” on page 1025
- “exportCertificate” on page 1026
- “exportCertToManagedKS” on page 1027
- “extractCertificate” on page 1028
- “getCertificate” on page 1029
- “getCertificateChain” on page 1029
- “importCertificate” on page 1030
- “importCertFromManagedKS” on page 1031
- “listPersonalCertificates” on page 1032
- “queryCACertificate” on page 1032
- “receiveCertificate” on page 1033
- “renewCertificate” on page 1034
- “replaceCertificate” on page 1035
- “requestCACertificate” on page 1036
- “revokeCACertificate” on page 1037

createChainedCertificate

The createChainedCertificate command creates a new self-signed certificate and stores the certificate in a keystore.

Note: To use the IBMi5OSKeyStore key store, verify that the signer for each part of the chain exists in the keystore before creating the new certificate. You must import the signer into the IBMi5OSKeyStore keystore before creating the new certificate.

Target object

None.

Required parameters

-keyStoreName

Specifies the name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

Specifies the name that uniquely identifies the certificate request in a keystore. (String, required)

-certificateSize

Specifies the size of the certificate. (Integer, required)

-certificateCommonName

Specifies the common name of the certificate. (String, required)

-certificateOrganization

Specifies the organization of the certificate. (String, optional)

Optional parameters

-rootCertificateAlias

Specifies a unique name to identify the root certificated to use for signing. The default root certificate alias is root. (String, optional)

-certificateVersion

Specifies the version of the certificate. (String, optional)

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-certificateOrganization

Specifies the organization of the certificate. (String, optional)

-certificateOrganizationalUnit

Specifies the organizational unit of the certificate. (String, optional)

-certificateLocality

Specifies the locality of the certificate. (String, optional)

-certificateState

Specifies the state of the certificate. (String, optional)

-certificateZip

Specifies the zip code of the certificate. (String, optional)

-certificateCountry

Specifies the country of the certificate. (String, optional)

-certificateValidDays

Specifies the amount of time in days for which the certificate is valid. (Integer, optional)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.createChainedCertificate('-keyStoreName myKeystore -certificateAlias  
newCertificate -certificateSize 10 -certificateCommonName localhost  
-certificateOrganization ibm')
```

- Using Jython list:

```
AdminTask.createChainedCertificate('-keyStoreName', 'myKeystore', '-certificateAlias',  
'newCertificate', '-certificateSize', '10', '-certificateCommonName', 'localhost',  
'-certificateOrganization', 'ibm')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createChainedCertificate('-interactive')
```

createSelfSignedCertificate

The createSelfSignedCertificate command creates a self-signed personal certificate in a keystore.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

-certificateVersion

The version of the certificate. (String, required)

-certificateSize

The size of the certificate. (Integer, required)

-certificateCommonName

The common name of the certificate. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

-certificateOrganization

The organization of the certificate. (String, optional)

-certificateOrganizationalUnit

The organizational unit of the certificate. (String, optional)

-certificateLocality

The locality of the certificate. (String, optional)

-certificateState

The state of the certificate. (String, optional)

-certificateZip

The zip code of the certificate. (String, optional)

-certificateCountry

The country of the certificate. (String, optional)

-certificateValidDays

The amount of time in days for which the certificate is valid. (Integer, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createSelfSignedCertificate {-keyStoreName testKeyStore -certificateAlias default -certificateCommonName localhost -certificateOrganization ibm}
```

- Using Jython string:

```
AdminTask.createSelfSignedCertificate(['-keyStoreName testKeyStore -certificateAlias default -certificateCommonName localhost -certificateOrganization ibm'])
```

- Using Jython list:

```
AdminTask.createSelfSignedCertificate(['-keyStoreName', 'testKeyStore', '-certificateAlias', 'default', '-certificateCommonName', 'localhost', '-certificateOrganization', 'ibm'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.createSelfSignedCertificate('-interactive')
```

deleteCertificate

The deleteCertificate command deletes a personal certificate from a keystore. The command saves a copy of the certificate in the delete keystore.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Interactive mode example usage:

- Using Jython:

```
AdminTask.deleteCertificate('-interactive')
```

exportCertificate

The exportCertificate command exports a personal certificate from one keystore to another.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-keyStorePassword

The password to the keystore. (String, required)

-keyFilePath

The full path to a keystore file that is located in a file system. The store from where a certificate will be imported or exported. (String, required)

-keyFilePassword

The password to the keystore file. (String, required)

-keyFileType

The type of the key file. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

-aliasInKeyStore
(String, optional)

Example output

The command does not return output.

Examples

Interactive mode example usage:

- Using Jython:

```
AdminTask.exportCertificate('-interactive')
```

exportCertToManagedKS

The `exportCertToManagedKS` command exports a personal certificate to a managed keystore in the configuration.

Target object

None.

Required parameters

-keyStoreName
Specifies the name that uniquely identifies the keystore configuration object. (String, required)

-keyStorePassword
The password to the keystore. (String, required)

-toKeyStoreName
Specifies the unique name of the keystore to export the certificate to. (String, required)

-certificateAlias
Specifies the alias of the certificate of interest. (String, required)

Optional parameters

-keyStoreScope
Specifies the keystore of the certificate of interest. (String, optional)

-toKeyStoreScope
Specifies the scope of the keystore to export to. (String, optional)

-aliasInKeyStore
Specifies the alias that identifies the certificate in the keystore. (String, optional)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.exportCertificateToManagedKS('-keyStoreName myKS -keyStorePassword myKSpw  
-toKeyStoreName myKS2 -certificateAlias testingKeyStore')
```

- Using Jython list:

```
AdminTask.exportCertificateToManagedKS(['-keyStoreName', 'myKS', '-keyStorePassword',  
'myKSpw', '-toKeyStoreName', 'myKS2', '-certificateAlias', 'testingKeyStore'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.exportCertificateToManagedKS('-interactive')
```

extractCertificate

The `extractCertificate` command extracts the signer part of a personal certificate to a certificate file. The certificate in the file can later be added to a keystore to establish trust.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

-certificateFilePath

The full path of the request file that contains the certificate. (String, required)

-base64Encoded

Set the value of this parameter to `true` if the certificate is a Base64 encoded ASCII file type. Set the value of this parameter to `false` if the certificate is binary. (Boolean, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl: Windows

```
$AdminTask extractCertificate {-keyStoreName testKeyStore -certificateFilePath
c:\temp\CertFile.arm -certificateAlias testCertificate}
```

AIX

HP-UX

AIX

Linux

Solaris

```
$AdminTask extractCertificate {-keyStoreName testKeyStore -certificateFilePath
/temp/CertFile.arm -certificateAlias testCertificate}
```

- Using Jython string: Windows

```
AdminTask.extractCertificate(['-keyStoreName testKeyStore -certificateFilePath
c:\temp\CertFile.arm -certificateAlias testCertificate'])
```

AIX

HP-UX

AIX

Linux

Solaris

```
AdminTask.extractCertificate(['-keyStoreName testKeyStore -certificateFilePath
/temp/CertFile.arm -certificateAlias testCertificate'])
```

- Using Jython list: Windows

```
AdminTask.extractCertificate(['-keyStoreName', 'testKeyStore', '-certificateFilePath',
'c:\temp\CertFile.arm', '-certificateAlias', 'testCertificate'])
```

AIX

HP-UX

AIX

Linux

Solaris

```
AdminTask.extractCertificate(['-keyStoreName', 'testKeyStore', '-certificateFilePath',
'/temp/CertFile.arm', '-certificateAlias', 'testCertificate'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.extractCertificate('-interactive')
```

getCertificate

The `getCertificate` command obtains information about a particular personal certificate in a keystore. If the certificate of interest was created with the `requestCACertificate` command, the certificate can be in the `COMPLETE` or `REVOKED` state. Certificate requests can be in the `PENDING` state. Use the `getCertificateRequest` command to determine if a certificate request is in the `PENDING` state.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command returns information about the certificate request.

Examples

Interactive mode example usage:

- Using Jython:

```
AdminTask.getCertificate('-interactive')
```

getCertificateChain

The `getCertificateChain` command queries your configuration for information about each personal certificate in a certificate chain.

Target object

None.

Required parameters and return values

-keyStoreName

Specifies the name of the keystore object that stores the CA certificate. Use the `listKeyStores` command to display a list of available keystores. (String, required)

-certificateAlias

Specifies the unique alias of the certificate. (String, required)

Optional parameters

-keyStoreScope

Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope. (String, optional)

Example output

The command returns an array of attribute lists that contain configuration information for each certificate in a chain.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getCertificateChain {-certificateAlias newCertificate  
-keyStoreName CellDefaultKeyStore}
```

- Using Jython string:

```
AdminTask.getCertificateChain('-certificateAlias newCertificate  
-keyStoreName CellDefaultKeyStore')
```

- Using Jython list:

```
AdminTask.getCertificateChain(['-certificateAlias', 'newCertificate',  
'-keyStoreName', 'CellDefaultKeyStore'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.getCertificateChain('-interactive')
```

importCertificate

The importCertificate command imports a personal certificate from a keystore.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-keyFilePath

The full path to a keystore file that is located in a file system. The store from where a certificate will be imported or exported. (String, required)

-keyFilePassword

The password to the keystore file. (String, required)

-keyFileType

The type of the key file. (String, required)

-certificateAliasFromKeyFile

The certificate alias in the key file from which the certificate is being imported. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Interactive mode example usage:

- Using Jython:

```
AdminTask.importCertificate('-interactive')
```

importCertFromManagedKS

The importCertFromManagedKS command imports a personal certificate from a managed keystore in the configuration.

Target object

None.

Required parameters

-keyStoreName

Specifies the name that uniquely identifies the keystore configuration object. (String, required)

-fromKeyStoreName

Specifies the name that uniquely identifies the keystore from which the system imports the certificate. (String, required)

-fromKeyStorePassword

Specifies the password for the keystore from which the system imports the certificate. (String, required)

-certificateAliasFromKeyStore

Specifies the alias of the certificate in the keystore. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope of the keystore to import the certificate to. (String, optional)

-fromKeyStoreScope

Specifies the scope of the keystore to import the certificate from. (String, optional)

-certificateAlias

Specifies the alias of the certificate for the destination keystore. (String, optional)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.importCertFromManagedKS('-keyStoreName myKeystore -fromKeyStoreName oldKeystore -fromKeyStorePassword my122password -certificateAliasFromKeyStore myCertificate')
```

- Using Jython list:

```
AdminTask.importCertFromManagedKS('-keyStoreName', 'myKeystore', '-fromKeyStoreName',
'oldKeystore', '-fromKeyStorePassword', 'my122password', '-certificateAliasFromKeyStore',
'myCertificate')
```

Interactive mode example usage

- Using Jython:

```
AdminTask.importCertFromManagedKS('-interactive')
```

listPersonalCertificates

The listPersonalCertificates command lists the personal certificates in a particular keystore.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command returns a list of attributes for each personal certificate in a keystore.

Examples

Batch mode example usage:

- Using Jython string:

```
AdminTask.listPersonalCertificates('-keyStoreName myKS')
```

- Using Jython list:

```
AdminTask.listPersonalCertificates(['-keyStoreName', 'myKS'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.listPersonalCertificates('-interactive')
```

queryCACertificate

The queryCACertificate command queries your configuration to determine if the CA has completed the certificate. If the CA returns a personal certificate, then the system marks the certificate as COMPLETE. Otherwise, it remains marked as PENDING.

Target object

None.

Required parameters and return values

-keyStoreName

Specifies the name of the keystore object that stores the CA certificate. Use the listKeyStores command to display a list of available keystores. (String, required)

-certificateAlias

Specifies the unique alias of the certificate. (String, required)

Optional parameters

-keyStoreScope

Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope. (String, optional)

Example output

The command returns one of two values: Certificate COMPLETE or certificate PENDING. If the command returns the Certificate COMPLETE message, the certificate authority returned the requested certificate and the default personal certificate is replaced. If the command returns the certificate PENDING message, the certificate authority did not yet return a certificate.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask queryCACertificate {-certificateAlias newCertificate
-keyStoreName CellDefaultKeyStore}
```

- Using Jython string:

```
AdminTask.queryCACertificate('-certificateAlias newCertificate
-keyStoreName CellDefaultKeyStore')
```

- Using Jython list:

```
AdminTask.queryCACertificate(['-certificateAlias', 'newCertificate',
'-keyStoreName', 'CellDefaultKeyStore'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.queryCACertificate('-interactive')
```

receiveCertificate

The receiveCertificate command receives a signer certificate from a file to a personal certificate.

Target object

None.

Required parameters

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

-certificateFilePath

The full path of the file that contains the certificate. (String, required)

-base64Encoded

Set the value of this parameter to true if the certificate is ascii base 64 encoded. Set the value of this parameter to false if the certificate is binary. (Boolean, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl: Windows

```
$AdminTask receiveCertificate {-keyStoreName testKeyStore  
-certificateFilePath c:\temp\CertFile.arm}
```

AIX

HP-UX

Linux

Solaris

```
$AdminTask receiveCertificate {-keyStoreName testKeyStore  
-certificateFilePath /temp/CertFile.arm}
```

- Using Jython string: Windows

```
AdminTask.receiveCertificate(['-keyStoreName testKeyStore  
-certificateFilePath c:\temp\CertFile.arm'])
```

AIX

HP-UX

Linux

Solaris

```
AdminTask.receiveCertificate(['-keyStoreName testKeyStore  
-certificateFilePath /temp/CertFile.arm'])
```

- Using Jython list: Windows

```
AdminTask.receiveCertificate(['-keyStoreName', 'testKeyStore',  
'-certificateFilePath', 'c:\temp\CertFile.arm'])
```

AIX

HP-UX

Linux

Solaris

```
AdminTask.receiveCertificate(['-keyStoreName', 'testKeyStore',  
'-certificateFilePath', '/temp/CertFile.arm'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.receiveCertificate('-interactive')
```

renewCertificate

The renewCertificate command renews a certificate with a new generated certificate.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name that identifies the keystore. (String, required)

-certificateAlias

Specifies the unique name that identifies the certificate. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope of the keystore. (String, optional)

-deleteOldSigners

Specifies whether to delete the old signers that are associated with the old certificate. Specify `false` to retain the old signers. (Boolean, optional)

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.renewCertificate(['-keyStoreName myKS -certificateAlias  
testCertificate'])
```

- Using Jython list:

```
AdminTask.renewCertificate(['-keyStoreName', 'myKS', '-certificateAlias',  
'testCertificate'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.renewCertificate('-interactive')
```

replaceCertificate

The `replaceCertificate` command replaces a personal certificate with another personal certificate. The command finds each reference to the old certificate alias in the configuration and replaces the alias with the new one. The command also replaces each signer certificate from the old personal certificate with the signer from the new personal certificate.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

The name that uniquely identifies the certificate request in a keystore. (String, required)

-replacementCertificateAlias

The alias of the certificate that is used to replace a different certificate. (String, required)

Optional parameters

-keyStoreScope

The scope name of the keystore. (String, optional)

-deleteOldCert

Set the value of this parameter to `true` if you want to delete the old signer certificates during certificate replacement. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

-deleteOldSigners

Set the value of this parameter to `true` if you want to delete the old certificates during certificate replacement. Otherwise, set the value of this parameter to `false`. (Boolean, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask replaceCertificate {-keyStoreName testKeyStore -certificateAlias
default -replacementCertificateAlias replaceCert -deleteOldCert true
-deleteOldSigners true}
```

- Using Jython string:

```
AdminTask.replaceCertificate(['-keyStoreName testKeyStore -certificateAlias
default -replacementCertificateAlias replaceCert -deleteOldCert true
-deleteOldSigners true'])
```

- Using Jython list:

```
AdminTask.replaceCertificate(['-keyStoreName', 'testKeyStore', '-certificateAlias',
'default', '-replacementCertificateAlias', 'replaceCert', '-deleteOldCert',
'true', '-deleteOldSigners', 'true'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.replaceCertificate('-interactive')
```

requestCACertificate

The `requestCACertificate` command creates a certificate request and sends the request to a certificate authority (CA). If the certificate authority returns a personal certificate, then the returned certificate replaces the certificate request in the keystore. The command also works with a preexisting certificate request that was created with the `createCertificateRequest` command. When the CA returns a personal certificate, the system marks the certificate as `COMPLETE` and the command returns a message stating that the certificate is complete. If the CA does not return a personal certificate, then the system marks the certificate request as `PENDING` and the command returns a message stating that the certificate is `PENDING`.

Note: To use the `IBMi5OSKeyStore` key store, verify that the signer for each part of the chain exists in the keystore before creating the new certificate. You must import the signer into the `IBMi5OSKeyStore` keystore before creating the new certificate.

Target object

None.

Required parameters and return values

-certificateAlias

Specifies the alias of the certificate. You can specify a predefined certificate request. (String, required)

-keyStoreName

Specifies the name of the keystore object that stores the CA certificate. Use the `listKeyStores` command to display a list of available keystores. (String, required)

-caClientName

Specifies the name of the CA client that was used to create the CA certificate. (String, required)

-revocationPassword

Specifies the password to use to revoke the certificate at a later date. (String, required)

Optional parameters

-keyStoreScope

Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope. (String, optional)

-caClientScope

Specifies the management scope of the CA client. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope. (String, optional)

-certificateCommonName

Specifies the common name (CN) part of the full distinguished name (DN) of the certificate. This common name can represent a person, company, or machine. For Web sites, the common name is frequently the DNS host name where the server resides. (String, optional)

-certificateOrganization

Specifies the organization part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateOrganizationalUnity

Specifies the organization unit part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateLocality

Specifies the locality part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateState

Specifies the state part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateZip

Specifies the zip code part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateCountry

Specifies the country part of the full distinguished name (DN) of the certificate. (String, optional)

-certificateSize

Specifies the size of the certificate key. The valid values are 512, 1024, and 2048. The default value is 1024. (String, optional)

Example output

The command returns one of two values: Certificate COMPLETE or certificate PENDING.

Examples**Batch mode example usage:**

- Using Jacl:

```
$AdminTask requestCACertificate {-certificateAlias newCertificate -keyStoreName
CellDefaultKeyStore -CAClientName myCAClient -revocationPassword revokeCApw}
```

- Using Jython string:

```
AdminTask.requestCACertificate('-certificateAlias newCertificate -keyStoreName
CellDefaultKeyStore -CAClientName myCAClient -revocationPassword revokeCApw')
```

- Using Jython list:

```
AdminTask.requestCACertificate(['-certificateAlias','newCertificate','-keyStoreName',
'CellDefaultKeyStore','-CAClientName','myCAClient','-revocationPassword',
'revokeCApw'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.requestCACertificate('-interactive')
```

revokeCACertificate

The revokeCACertificate command sends a request to the CA to revoke the CA personal certificate of interest.

Target object

None.

Required parameters and return values

-certificateAlias

Specifies the unique name that identifies the CA personal certificate object and the alias name of the certificate in the keystore. (String, required)

-keyStoreName

Specifies the name of the keystore where the CA personal certificate is stored. (String, required)

-revocationPassword

Specifies the password needed to revoke the certificate. This is the same password that was provided when the certificate was created. (String, required)

Optional parameters

-keyStoreScope

Specifies the management scope of the keystore. For a deployment manager profile, the default value is the cell scope. For an application server profile, the default value is the node scope. (String, optional)

-revocationReason

Specifies the reason for revoking the certificate of interest. The default value for this parameter is unspecified. (String, optional)

Example output

The command does not return output. Use the `getCertificate` command to view the current status of the certificate, as the following example displays:

```
AdminTask.getCertificate('-certificateAlias myCertificate -keyStoreName CellDefaultKeyStore')
```

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask revokeCACertificate {-keyStoreName CellDefaultKeyStore -certificateAlias myCertificate -revocationPassword pw4revoke}
```

- Using Jython string:

```
AdminTask.revokeCACertificate(['-keyStoreName CellDefaultKeyStore -certificateAlias myCertificate -revocationPassword pw4revoke'])
```

- Using Jython list:

```
AdminTask.revokeCACertificate(['-keyStoreName', 'CellDefaultKeyStore', '-certificateAlias', 'myCertificate', '-revocationPassword', 'pw4revoke'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.revokeCACertificate('-interactive')
```

WSCertExpMonitorCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the `wsadmin` tool. The commands and parameters in the `WSCertExpMonitorCommands` group can be used to start or update the certificate expiration monitor.

The `WSCertExpMonitorCommands` command group for the `AdminTask` object includes the following commands:

- “createWSCertExpMonitor”
- “deleteWSCertExpMonitor” on page 1040
- “getWSCertExpMonitor” on page 1040
- “listWSCertExpMonitor” on page 1041
- “modifyWSCertExpMonitor” on page 1042
- “startCertificateExpMonitor” on page 1043

createWSCertExpMonitor

The **createWSCertExpMonitor** command creates the certificate expiration monitor settings in the configuration.

Target object

None.

Required parameters and return values

-name

The name that uniquely identifies the certificate expiration monitor. (String, required)

-autoReplace

Set the value of this parameter to true if you want to replace a certificate within a certificate expiration date. If not, set the value of this parameter to false. (Boolean, required)

-deleteOld

Set the value of this parameter to true if you want to delete an old certificate during certificate expiration monitoring. If not, set the value of this parameter to false. (Boolean, required)

-daysBeforeNotification

The number of days before a certificate expires that you want to be notified of the expiration. (Integer, required)

-wsScheduleName

The name of the scheduler to use for certificate expiration. (String, required)

-wsNotificationName

The name of the notifier to use for certificate expiration. (String, required)

-isEnabled

Set the value of this parameter to true if the certificate expiration monitor is enabled. If not, set the value of this parameter to false. (Boolean, optional)

- Returns: The configuration object name of the certificate expiration monitor object that you created.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createWSCertExpMonitor {-name testCertMon -autoReplace true -deleteOld true
-daysBeforeNotification 30 -wsScheduleName testSchedule -wsNotificationName testNotifier
-isEnabled false}
```

- Using Jython string:

```
AdminTask.createWSCertExpMonitor ('[-name testCertMon -autoReplace true -deleteOld true
-daysBeforeNotification 30 -wsScheduleName testSchedule -wsNotificationName testNotifier
-isEnabled false]')
```

- Using Jython list:

```
AdminTask.createWSCertExpMonitor (['-name', 'testCertMon', '-autoReplace', 'true', '-deleteOld',
'true', '-daysBeforeNotification', '30', '-wsScheduleName', 'testSchedule', '-wsNotificationName',
'testNotifier', '-isEnabled', 'false'])
```

Interactive mode example usage:

- Using Jacl:
`$AdminTask createWSCertExpMonitor {-interactive}`
- Using Jython string:
`AdminTask.createWSCertExpMonitor ('[-interactive]')`
- Using Jython list:
`AdminTask.createWSCertExpMonitor (['-interactive'])`

deleteWSCertExpMonitor

The **deleteWSCertExpMonitor** command deletes the settings of a scheduler from the configuration.

Target object

None.

Required parameters and return values

-name

The name that uniquely identifies the certificate expiration monitor. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteWSCertExpMonitor {-name testCertMon}`
- Using Jython string:
`AdminTask.deleteWSCertExpMonitor ('[-name testCertMon]')`
- Using Jython list:
`AdminTask.deleteWSCertExpMonitor (['-name', 'testCertMon'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteWSCertExpMonitor {-interactive}`
- Using Jython string:
`AdminTask.deleteWSCertExpMonitor ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteWSCertExpMonitor (['-interactive'])`

getWSCertExpMonitor

The **getWSCertExpMonitor** command displays the settings of a particular scheduler.

Target object

None.

Required parameters and return values

-name

The name that uniquely identifies the certificate expiration monitor. (String, required)

- Returns: The scheduler in the configuration.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getWSCertExpMonitor {-name testCertMon}
```
- Using Jython string:

```
AdminTask getWSCertExpMonitor ('[-name testCertMon]')
```
- Using Jython list:

```
AdminTask getWSCertExpMonitor (['-name', 'testCertMon'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getWSCertExpMonitor {-interactive}
```
- Using Jython string:

```
AdminTask.getWSCertExpMonitor ('[-interactive]')
```
- Using Jython list:

```
AdminTask.getWSCertExpMonitor (['-interactive'])
```

listWSCertExpMonitor

The **listWSCertExpMonitor** command lists the scheduler in the configuration.

Target object

None.

Required parameters and return values

-displayObjectNames

If you set the value of this parameter to true, the command returns the certificate expiration monitor configuration object. If you set the value of this parameter to false, the command returns the name of the certificate expiration monitor. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listWSCertExpMonitor {-displayObjectName false}
```
- Using Jython string:

```
AdminTask.listWSCertExpMonitor ('[-displayObjectName false]')
```
- Using Jython list:

```
AdminTask.listWSCertExpMonitor (['-displayObjectName', 'false'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listWSCertExpMonitor {-interactive}
```
- Using Jython string:

```
AdminTask.listWSCertExpMonitor ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listWSCertExpMonitor (['-interactive'])
```

modifyWSCertExpMonitor

The **modifyWSCertExpMonitor** command changes the setting of an existing scheduler.

Target object

None.

Required parameters and return values

-name

The name that uniquely identifies the certificate expiration monitor. (String, required)

-autoReplace

Set the value of this parameter to true if you want to replace a certificate within a certificate expiration date. If not, set the value of this parameter to false. (Boolean, required)

-deleteOld

Set the value of this parameter to true if you want to delete an old certificate during certificate expiration monitoring. If not, set the value of this parameter to false. (Boolean, required)

-daysBeforeNotification

The number of days before a certificate expires that you want to be notified of the expiration. (Integer, required)

-wsScheduleName

The name of the scheduler to use for certificate expiration. (String, required)

-wsNotificationName

The name of the notifier to use for certificate expiration. (String, required)

-isEnabled

Set the value of this parameter to true if the certificate expiration monitor is enabled. If not, set the value of this parameter to false. (Boolean, optional)

- Returns: None

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifyWSCertExpMonitor {-name testCertMon -autoReplace false -deleteOld false -daysBeforeNotification 20 -isEnabled true}
```

- Using Jython string:

```
AdminTask.modifyWSCertExpMonitor (['-name testCertMon -autoReplace false -deleteOld false -daysBeforeNotification 20 -isEnabled true'])
```

- Using Jython list:

```
AdminTask.modifyWSCertExpMonitor (['-name', 'testCertMon', '-autoReplace', 'false', '-deleteOld', 'false', '-daysBeforeNotification', '20', '-isEnabled', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifyWSCertExpMonitor {-interactive}
```

- Using Jython string:

```
AdminTask.modifyWSCertExpMonitor (['-interactive'])
```

- Using Jython list:
`AdminTask.modifyWSCertExpMonitor (['-interactive'])`

startCertificateExpMonitor

The **startCertificateExpMonitor** command performs certificate monitoring. This command visits all key stores and checks to see if they are within certificate expiration range.

Target object

None.

Required parameters and return values

- Parameters: None
- Returns: None

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask startCertificateExpMonitor`
- Using Jython:
`AdminTask.startCertificateExpMonitor()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask startCertificateExpMonitor {-interactive}`
- Using Jython string:
`AdminTask.startCertificateExpMonitor (['-interactive'])`
- Using Jython list:
`AdminTask.startCertificateExpMonitor (['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

SignerCertificateCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the SignerCertificateCommands group can be used to create and modify signer certificates in relation to the key store file and to query for signer information on ports of remote hosts.

The SignerCertificateCommands command group for the AdminTask object includes the following commands:

- “addSignerCertificate” on page 1044
- “deleteSignerCertificate” on page 1045
- “extractSignerCertificate” on page 1045
- “getSignerCertificate” on page 1046

- “listSignerCertificates” on page 1047
- “retrieveSignerFromPort” on page 1047
- “retrieveSignerInfoFromPort” on page 1048

addSignerCertificate

The **addSignerCertificate** command add a signer certificate from a certificate file to a keystore.

Target object

None.

Required parameters

-keyStoreName

Specifies the name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

Specifies the name that uniquely identifies the certificate request in a keystore. (String, required)

-certificateFilePath

Specifies the full path of the request file that contains the certificate. (String, required)

-base64Encoded

Specifies that the certificate is a Base64 encoded ASCII data file type if the value is set to true. Set the value of this parameter to false if the certificate is a binary DER data file type. (Boolean, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask addSignerCertificate {-keyStoreName testKeyStore -certificateAlias
default -certificateFilePath <file path> -base64Encoded true}
```
- Using Jython string:


```
AdminTask.addSignerCertificate(['-keyStoreName testKeyStore -certificateAlias
default -certificateFilePath <file path> -base64Encoded true'])
```
- Using Jython list:


```
AdminTask.addSignerCertificate(['-keyStoreName', 'testKeyStore', '-certificateAlias',
'default', '-certificateFilePath', '<file path>', '-base64Encoded', 'true'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask addSignerCertificate {-interactive}
```
- Using Jython string:


```
AdminTask.addSignerCertificate ('[-interactive]')
```

deleteSignerCertificate

The **deleteSignerCertificate** command delete a signer certificate from a certificate file from a keystore.

Target object

None.

Required parameters

-keyStoreName

Specifies the name that uniquely identifies the keystore configuration object. (String, required)

-certificateAlias

Specifies the name that uniquely identifies the certificate request in a keystore. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

Example output

The command does not return output.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteSignerCertificate {-keyStoreName testKeyStore -certificateAlias default}
```

- Using Jython string:

```
AdminTask.deleteSignerCertificate(['-keyStoreName testKeyStore -certificateAlias default'])
```

- Using Jython list:

```
AdminTask.deleteSignerCertificate(['-keyStoreName', 'testKeyStore', '-certificateAlias', 'default'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteSignerCertificate {-interactive}
```

- Using Jython string:

```
AdminTask.deleteSignerCertificate (['-interactive'])
```

extractSignerCertificate

The **extractSignerCertificate** command extracts a signer certificate from a key store to a file.

Target object

None

Required parameters and return values

-keyStoreName

The name of the key store where the signer certificate is located. (String, required)

-keyStoreScope

The management scope of the key store. (String, optional)

-certificateAlias

The alias name of the signer certificate in the key store. (String, required)

-certificateFilePath

The full path name of the file that contains the signer certificate. (String, required)

-base64Encoded

Set the value of this parameter to true if the certificate is ascii base 64 encoded. Set the value of this parameter to false if the certificate is binary. (String, required)

Examples**Interactive mode example usage:**

- Using Jacl:
`$AdminTask extractSigner Certificate {-interactive}`
- Using Jython string:
`AdminTask.extractSignerCertificate ('[-interactive]')`
- Using Jython list:
`AdminTask.extractSignerCertificate (['-interactive'])`

getSignerCertificate

The **getSignerCertificate** command obtains information about a signer certificate from a key store.

Target object

None

Required parameters and return values

-keyStoreName

The name of the key store where the signer certificate is located. (String, required)

-keyStoreScope

The management scope of the key store. (String, optional)

-certificateAlias

The alias name of the signer certificate in the key store. (String, required)

Examples**Interactive mode example usage:**

- Using Jacl:
`$AdminTask getSignerCertificate {-interactive}`
- Using Jython string:
`AdminTask.getSignerCertificate ('[-interactive]')`
- Using Jython list:
`AdminTask.getSignerCertificate (['-interactive'])`

listSignerCertificates

The **listSignerCertificates** command lists all signer certificates in a particular key store.

Target object

None

Required parameters and return values

-keyStoreName

The name of the key store where the signer certificate is located. (String, required)

-keyStoreScope

The management scope of the key store. (String, optional)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask listSignerCertificates {-interactive}`
- Using Jython string:
`AdminTask.listSignerCertificates ('[-interactive]')`
- Using Jython list:
`AdminTask.listSignerCertificates (['-interactive'])`

retrieveSignerFromPort

The **retrieveSignerFromPort** command retrieves a signer from a remote host and stores the signer in a key store.

Target object

None

Required parameters and return values

-host

The host name of the system from where the signer certificate will be retrieved. (String, required)

-port

The port of the remote system from where the signer certificate will be retrieved. (Integer, required)

-keyStoreName

The name of the key store where the signer certificate is located. (String, required)

-keyStoreScope

The management scope of the key store. (String, required)

-sslConfigName

The name of the SSL configuration object. (String, optional)

-sslConfigScopeName

The management scope where the SSL configuration object is located. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask retrieveSigner FromPort {-host serverHost -port 443 -keyStoreName testKeyStore
-certificate Alias serverHostSigner}
```
- Using Jython string:


```
AdminTask.retrieveSigner FromPort (['-host server Host -port 443 -keyStore Name testKeyStore
-certificateAlias serverHost Signer']')
```
- Using Jython list:


```
AdminTask.retrieveSigner FromPort (['-host', 'serverHost', '-port', '443', '-keyStoreName',
'testKeyStore', '-certificateAlias', 'serverHost Signer'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask retrieveSigner FromPort {-interactive}
```
- Using Jython string:


```
AdminTask.retrieveSigner FromPort (['-interactive']')
```
- Using Jython list:


```
AdminTask.retrieveSigner FromPort (['-interactive'])
```

retrieveSignerInfoFromPort

The **retrieveSigner InfoFromPort** command retrieves signer information from a port on a remote host.

Target object

None

Required parameters and return values

-host

The host name of the system from where the signer certificate will be retrieved. (String, required)

-port

The port of the remote system from where the signer certificate will be retrieved. (Integer, required)

-sslConfigName

The name of the SSL configuration object. (String, optional)

-sslConfigScopeName

The management scope where the SSL configuration object is located. (String, optional)

Examples

Interactive mode example usage:

- Using Jacl:


```
$AdminTask retrieveSigner InfoFromPort {-interactive}
```
- Using Jython string:


```
AdminTask.retrieveSigner InfoFromPort (['-interactive']')
```
- Using Jython list:


```
AdminTask.retrieveSigner InfoFromPort (['-interactive'])
```

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

CertificateRequestCommands command group of the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the CertificateRequestCommands group can be used to create and manage certificate requests.

The CertificateRequestCommands command group for the AdminTask object includes the following commands:

- “createCertificateRequest”
- “deleteCertificateRequest” on page 1050
- “extractCertificateRequest” on page 1051
- “getCertificateRequest” on page 1051
- “listCertificateRequest” on page 1052

createCertificateRequest

The **createCertificateRequest** command creates a certificate request that is associated with a particular key store.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the key store configuration object. (String, required)

-keyStoreScope

The scope name of the key store. (String, optional)

-certificateAlias

The name that uniquely identifies the certificate request in a key store. (String, required)

-certificateVersion

The certificate version. (String, required)

-certificateSize

(Integer, required)

-certificateCommonName

(String, required)

-certificateOrganization

(String, optional)

-certificateOrganizationalUnit

(String, optional)

-certificateLocality

(String, optional)

-certificateState

The state code for the certificate. (String, optional)

-certificateZip

The zip code for the certificate. (String, optional)

-certificateCountry

The country for the certificate. (String, optional)

-certificateValidDays

The amount of time in days for which the certificate is valid. (Integer, optional)

-certificateRequestFilePath

The file location of the certificate request that can be sent to a certificate authority. (String, required)

- Returns: The configuration object name of the key store object that you created.

Examples**Batch mode example usage:**

- Using Jacl:

```
$AdminTask createCertificateRequest {-keyStoreName testKeyStore  
-certificateAlias certReq -certificateSize 1024 -certificate  
CommonName localhost -certificate Organization testing -certificate  
RequestFilePath c:\temp\testCertReq.arm}
```

- Using Jython string:

```
AdminTask.createCertificateRequest ('[-keyStoreName testKeyStore  
-certificateAlias certReq -certificateSize 1024 -certificate  
CommonName localhost -certificate Organization testing -certificate  
RequestFilePath c:\temp\testCertReq.arm]')
```

- Using Jython list:

```
AdminTask.createCertificateRequest (['-keyStoreName', 'testKeyStore',  
'-certificateAlias', 'certReq', '-certificateSize', '1024',  
'-certificateCommonName', 'localhost', '-certificateOrganization',  
'testing', '-certificateRequestFilePath', 'c:\temp\testCertReq.arm'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createCertificateRequest {-interactive}
```

- Using Jython string:

```
AdminTask.createCertificateRequest ('[-interactive]')
```

- Using Jython list:

```
AdminTask.createCertificateRequest (['-interactive'])
```

deleteCertificateRequest

The **deleteCertificateRequest** command deletes a certificate request from a key store.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the key store configuration object. (String, required)

-keyStoreScope

The scope name of the key store. (String, optional)

-certificateAlias

The name that uniquely identifies the certificate request in a key store. (String, required)

- Returns: None.

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteCertificateRequest {-interactive}`
- Using Jython string:
`AdminTask.deleteCertificateRequest ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteCertificateRequest (['-interactive'])`

extractCertificateRequest

The **extractCertificateRequest** command extracts a certificate request to a file.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the key store configuration object. (String, required)

-keyStoreScope

The scope name of the key store. (String, optional)

-certificateAlias

The name that uniquely identifies the certificate request in a key store. (String, required)

-certificateRequestFilePath

The file location of the certificate request that can be sent to a certificate authority. (String, required)

- Returns: A certificate request file is created that contains the extracted certificate.

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask extractCertificateRequest {-interactive}`
- Using Jython string:
`AdminTask.extractCertificateRequest ('[-interactive]')`
- Using Jython list:
`AdminTask.extractCertificateRequest (['-interactive'])`

getCertificateRequest

The **getCertificateRequest** command obtains information about a particular certificate request in a key store.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the key store configuration object. (String, required)

-keyStoreScope

The scope name of the key store. (String, optional)

-certificateAlias

The name that uniquely identifies the certificate request in a key store. (String, required)

- Returns: Information about the certificate request.

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask getCertificateRequest {-interactive}`
- Using Jython string:
`AdminTask.getCertificateRequest ('[-interactive]')`
- Using Jython list:
`AdminTask.getCertificateRequest (['-interactive'])`

listCertificateRequest

The **listCertificateRequest** command lists all the certificate requests associated with a particular key store.

Target object

None.

Required parameters and return values

-keyStoreName

The name that uniquely identifies the key store configuration object. (String, required)

-keyStoreScope

The scope name of the key store. (String, optional)

- Returns: An attribute list for each certificate request in a key store.

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask listCertificateRequest {-interactive}`
- Using Jython string:
`AdminTask.listCertificateRequest ('[-interactive]')`
- Using Jython list:
`AdminTask.listCertificateRequest (['-interactive'])`

Enabling authentication in the file transfer service using scripting

You can enable authentication in the file transfer service using scripting and the wsadmin tool.

Before you begin

Before starting this task, the wsadmin tool must be running. See the Starting the wsadmin scripting client article for more information.

About this task

In WebSphere Application Server Network Deployment, V5.0.1 or later, the file transfer service is enhanced to provide 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, V5.0 and V5.0.1 or later.

Turning the file transfer authentication on is recommended to prevent unauthorized use of the file transfer application; however, if you have any V5.0 clients in your Network Deployment environment, they cannot communicate with the secured file transfer application if global security is turned on.

In WebSphere Application Server V6.x, mixed cells are supported and file transfer has become a system application. If all of the nodes in the cell are of V5.0.1 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.

- A wsadmin Jacl script is provided to help you redeploy the file transfer. The script is called *redeployFileTransfer.jacl* and is located in the *app_server_root/bin* directory.

After the deployment manager and all the nodes are upgraded to WebSphere Application Server Network Deployment, 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 is the following:

–

```
wsadmin -conntype NONE -lang jacl -profile redeployFileTransfer.jacl -c "fileTransferAuthenticationXxx  
cellName nodeName serverName"
```

where Xxx is **On** or **Off**.

Windows Use wsadmin or wsadmin.bat.

Linux **UNIX** Use wsadmin.sh.

- For example, when running the script to enable use of the *filetransferSecured.ear* file, the syntax is similar to the following example:

```
wsadmin -conntype NONE -lang jacl -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOn managedCell managedCellManager dmgr"
```

or

```
wsadmin -conntype NONE -lang jacl -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOn baseCell base server1"
```

- If you want to go return to running the file transfer service without authentication, you can run the script as shown in the following example:

```
wsadmin -conntype NONE -lang jacl -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOff baseNodeCell baseNode server1"
```

or

```
wsadmin -conntype NONE -lang jacl -profile redeployFileTransfer.jacl -c
"fileTransferAuthenticationOff managedCell managedCellManager dmgr"
```

What to do next

You must restart the server for the change to take affect.

Propagating security policy of installed applications to a JACC provider using wsadmin scripting

It is possible that you have applications installed prior to enabling the Java Authorization Contract for Containers (JACC)-based authorization. You can start with default authorization and then move to an external provider-based authorization using JACC later.

Before you begin

Note: Use the wsadmin tool to propagate information to the JACC provider independent of the application installation process, avoiding the need to reinstall applications. Also, during application installation or modification you might have had problems propagating the security policy information to the JACC provider. For example, network problems might occur, the JACC provider might not be available, and so on. For these cases, the security policy of the previously installed applications does not exist in the JACC provider to make the access decisions. One choice is to reinstall the applications involved. However, you can avoid reinstalling by using the wsadmin scripting tool. Use this tool to propagate information to the JACC provider independent of the application installation process. The tool eliminates the need for reinstalling the applications.

The tool uses the SecurityAdmin MBean to propagate the policy information in the deployment descriptor of any installed application to the JACC provider. You can invoke this tool using wsadmin at the base application server for base and deployment manager level for Network Deployment. Note that the SecurityAdmin MBean is available only when the server is running.

Use `propagatePolicyToJACCProvider{-appNames appNames}` to propagate the policy information in the deployment descriptor or annotations of the enterprise archive (EAR) files to the JACC provider. If the `RoleConfigurationFactory` and the `RoleConfiguration` interfaces are implemented by the JACC provider, the authorization table information in the binding file of the EAR files is also propagated to the provider. See the *Securing applications and their environment* PDF for more information about these interfaces.

The `appNames` String contains the list of application names, delimited by a colon (:), whose policy information must be stored in the provider. If `appNames` is not present, the policy information of all the deployed applications is propagated to the provider.

Also, be aware of the following items:

- Before migrating applications to the Tivoli Access Manager JACC provider, create or import the users and groups that are in the applications to Tivoli Access Manager.
- Depending on the application or the number of applications that are propagated, you might have to increase the request time-out period either in the `soap.client.props` file in the directory `profile_root/properties` (if using SOAP) or in the `sas.client.props` file (if using RMI) for the command to complete. You can set the request time-out value to 0 to avoid the timeout problem, and change it back to the original value after the command is run.

1. Configure your JACC provider in WebSphere Application Server.
See the *Securing applications and their environment* PDF for more information.
2. Restart the server.
3. Enter the following commands:

```
wsadmin>$AdminTask propagatePolicyToJACCProvider {-appNames appNames}
```


JACCUtilityCommands command group for the AdminTask object

Use this topic as a reference for the commands for the JACCUtilityCommands group for the AdminTask object. Use these commands to determine whether Java Authorization Contract for Containers (JACC) is enabled and whether the runtime uses a single security domain. You can also use these commands to propagate the security policies for application to the JACC provider.

The following commands are available for the JACCUtilityCommands group of the AdminTask object.

- “isJACCEnabled”
- “isSingleSecurityDomain”
- “propagatePolicyToJACCProvider ” on page 1056

isJACCEnabled

The isJACCEnabled command displays whether JACC is enabled or disabled in the global security domain when the server was started. The command does not indicate dynamic changes. Instead, it displays the JACC status at server startup.

Target object

None.

Required parameters

None.

Return value

The command returns true if JACC is enabled. The command returns false if JACC is disabled.

Batch mode example usage

Using Jython string:

```
AdminTask.isJACCEnabled()
```

Interactive mode example usage

Using Jython:

```
AdminTask.isJACCEnabled('-interactive')
```

isSingleSecurityDomain

The isSingleSecurityDomain command displays whether the environment is configured to use a single security domain when the server was started. The command does not indicate dynamic changes. Instead, it displays the security domain status at server startup.

Target object

None.

Required parameters

None.

Return value

The command returns true if the environment uses a single security domain. The command returns the false string if the environment uses multiple security domains.

Batch mode example usage

Using Jython:

```
AdminTask.isSingleSecurityDomain()
```

Interactive mode example usage

Using Jython:

```
AdminTask.isSingleSecurityDomain('-interactive')
```

propagatePolicyToJACCPProvider

The propagatePolicyToJACCPProvider command propagates the security policies of the applications of interest to the JACC provider. This command is supported in a single security domain environment only.

Target object

None.

Required parameters

None.

Optional parameters

-appNames

Specifies a list of application names delimited with a colon character (:). (String, optional)

The command uses all applications if you do not specify a value for this parameter, as the following syntax demonstrates: `AdminTask.propagatePolicyToJACCPProvider()`

Return value

The command does not return output.

Batch mode example usage

Using Jython string:

```
AdminTask.propagatePolicyToJACCPProvider ('-appNames "app1:app2:app3"')
```

Using Jython list:

```
AdminTask.propagatePolicyToJACCPProvider ('-appNames', ' "app1:app2:app3"')
```

Interactive mode example usage

Using Jython:

```
AdminTask.propagatePolicyToJACCPProvider ('-interactive')
```

Configuring custom adapters for federated repositories using wsadmin

You can use the Jython or Jacl scripting language with the wsadmin tool to define custom adapters in the federated repositories configuration file.

Before you begin

Shut down the WebSphere Application Server and the wsadmin command window.

About this task

The federated repositories configuration file, `wimconfig.xml`, is shipped with WebSphere Application Server 6.1.x and is located in the `app_server_root/profiles/profile_name/config/cells/cell_name/wim/config` directory.

Note: For additional information about the commands to use for this topic, see “`IdMgrRepositoryConfig` command group for the `AdminTask` object” on page 1163.

Use the following steps to add a custom adapter to any federated repositories configuration file and to any realm defined within the configuration file.

1. Open the `wimconfig.xml` file with a text editor.
2. Add a new `config:repositories` element to the file. This element should be placed before the `config:realmConfiguration` element.

The following example configures a custom repository to use the `com.ibm.ws.wim.adapter.sample.SampleFileAdapter` class and sets the `SampleFileRepository` repository as the identifier:

```
<config:repositories adapterClassName="com.ibm.ws.wim.adapter.sample.SampleFileAdapter"
id="SampleFileRepository"/>
```

3. Save the `wimconfig.xml` file and close the text editor.
4. Copy the `vmmsampleadapter.jar` file that is provided to `app_server_root/lib`.
5. Start the wsadmin command prompt application and enter the following command:

```
wsadmin -conntype none
```

6. Disable paging in the common repository configuration. Set the `supportPaging` parameter for the `updateIdMgrRepository` command to `false` to disable paging.

Note: You must perform this step because the sample adapter does not support paging.

The following examples use the `SampleFileRepository` repository as the identifier for the custom repository.

Using Jython:

```
AdminTask.updateIdMgrRepository('-id SampleFileRepository -supportPaging false')
```

Using Jacl:

```
$AdminTask updateIdMgrRepository {-id SampleFileRepository -supportPaging
false}
```

Note: A warning will appear until the configuration of the sample repository is complete.

7. Add the necessary custom properties for the adapter. Use the `setIdMgrCustomProperty` command repeatedly to add multiple properties. Use this command once per property to add multiple properties to your configuration. You must use both the `name` and `value` parameters to add the custom property for the specified repository. For example, to add a custom property of `fileName`, enter the following command.

Using Jython:

```
AdminTask.setIdMgrCustomProperty('-id SampleFileRepository -name fileName
-value "c:\sampleFileRegistry.xml"')
```

Using Jacl:

```
$AdminTask setIdMgrCustomProperty {-id SampleFileRepository -name fileName
-value "c:\sampleFileRegistry.xml"}
```

8. Add a base entry to the adapter configuration. Use the `addIdMgrRepositoryBaseEntry` command to specify the name of the base entry for the specified repository. For example:

Using Jython:

```
AdminTask.addIdMgrRepositoryBaseEntry('-id SampleFileRepository -name  
o=sampleFileRepository')
```

Using Jacl:

```
$AdminTask addIdMgrRepositoryBaseEntry {-id SampleFileRepository -name  
o=sampleFileRepository}
```

9. Use the `addIdMgrRealmBaseEntry` command to add the base entry to the realm, which will link the realm with the repository:

Using Jython:

```
AdminTask.addIdMgrRealmBaseEntry('-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository')
```

Using Jacl:

```
$AdminTask addIdMgrRealmBaseEntry {-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository}
```

10. Save your configuration changes. Enter the following commands to save the new configuration and close the `wsadmin` tool.

Using Jython:

```
AdminConfig.save()  
exit
```

Using Jacl:

```
$AdminConfig save  
exit
```

The following example displays the complete text of the newly-revised `wimconfig.xml` file:

```
<!--  
  Begin Copyright  
  
  Licensed Materials - Property of IBM  
  
  virtual member manager  
  
  (C) Copyright IBM Corp. 2005 All Rights Reserved.  
  
  US Government Users Restricted Rights - Use, duplication or  
  disclosure restricted by GSA ADP Schedule Contract with IBM Corp.  
  
  End Copyright  
-->  
<sdo:datagraph xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
xmlns:config="http://www.ibm.com/websphere/wim  
/config" xmlns:sdo="commonj.sdo">  
  <config:configurationProvider maxPagingResults="500" maxSearchResults="4500"  
maxTotalPagingResults="1000"  
pagedCacheTimeOut="900" pagingEntityObject="true" searchTimeOut="600000">  
    <config:dynamicModel xsdFileName="wimdatagraph.xsd"/>  
    <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="Group">  
      <config:rdnProperties>cn</config:rdnProperties>  
    </config:supportedEntityTypes>  
    <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="OrgContainer">  
      <config:rdnProperties>o</config:rdnProperties>  
      <config:rdnProperties>ou</config:rdnProperties>  
      <config:rdnProperties>dc</config:rdnProperties>  
      <config:rdnProperties>cn</config:rdnProperties>  
    </config:supportedEntityTypes>  
    <config:supportedEntityTypes defaultParent="o=defaultWIMFileBasedRealm" name="PersonAccount">  
      <config:rdnProperties>uid</config:rdnProperties>  
    </config:supportedEntityTypes>  
    <config:repositories xsi:type="config:FileRepositoryType" adapterClassName="com.ibm.  
ws.wim.adapter.file.was.FileAdapter"  
id="InternalFileRepository" supportPaging="false" supportSorting="false" messageDigestAlgorithm="SHA-1">
```

```

    <config:baseEntries name="o=defaultWIMFileBasedRealm"/>
</config:repositories>
<config:repositories adapterClassName="com.ibm.ws.wim.adapter.sample.SampleFileAdapter"
id="SampleFileRepository">
    <config:CustomProperties name="fileName" value="c:\sampleFileRegistry.xml"/>
    <config:baseEntries name="o=sampleFileRepository"/>
</config:repositories>
<config:realmConfiguration defaultRealm="defaultWIMFileBasedRealm">
    <config:realms delimiter="@" name="defaultWIMFileBasedRealm" securityUse="active">
        <config:participatingBaseEntries name="o=defaultWIMFileBasedRealm"/>
        <config:participatingBaseEntries name="o=sampleFileRepository"/>
        <config:uniqueUserIdMapping propertyForInput="uniqueName" propertyForOutput="uniqueName"/>
        <config:userSecurityNameMapping propertyForInput="principalName" propertyForOutput="principalName"/>
        <config:userDisplayNameMapping propertyForInput="principalName" propertyForOutput="principalName"/>
        <config:uniqueGroupIdMapping propertyForInput="uniqueName" propertyForOutput="uniqueName"/>
        <config:groupSecurityNameMapping propertyForInput="cn" propertyForOutput="cn"/>
        <config:groupDisplayNameMapping propertyForInput="cn" propertyForOutput="cn"/>
    </config:realms>
</config:realmConfiguration>
</config:configurationProvider></sdo:datagraph>

```

11. Restart the application server.

Related reference

“Sample custom adapters for federated repositories examples” on page 223

Out of the box adapters for federated repositories provide File, LDAP, and Database adapters for your use. These adapters implement the `com.ibm.wsspi.wim.Repository` software programming interface (SPI). A virtual member manager custom adapter needs to implement the same SPI.

Disabling embedded Tivoli Access Manager client using wsadmin

Follow these steps to unconfigure the Java Authorization Contract for Containers (JACC) provider for Tivoli Access Manager.

About this task

In a Network Deployment architecture, ensure that all the managed servers, including node agents, are started. Perform the following process once on the deployment management server. Details of the unconfiguration are forwarded to managed servers, including node agents, when a synchronization is performed. The managed servers require their own reboot for the configuration changes to take effect.

Note: It is also possible to unconfigure using the administrative console. For details on unconfiguring the embedded Tivoli Access Manager client using the WebSphere Application Server administrative console, refer to “Disabling embedded Tivoli Access Manager client using the administrative console” on page 559.

1. Restart the deployment manager process.
2. Start the `wsadmin` command-line utility. The `wsadmin` command is found in the `install_dir/bin` directory
3. From the **wsadmin** prompt, enter the following command:

```
WSADMIN>$AdminTask unconfigureTAM -interactive
```

You are prompted to enter the following information:

Option	Description
WebSphere Application Server node name	Enter an asterisk (*) to select all nodes.
Tivoli Access Manager administrator user name	Enter the Tivoli Access Manager administration user ID, as created at the time of Tivoli Access Manager configuration. This name is usually, <code>sec_master</code> .

Option	Description
Tivoli Access Manager administrator user password	Enter the password for the Tivoli Access Manager administrator.
Force	Enter <i>yes</i> , if you want to ignore errors when unconfiguring the JACC provider for Tivoli Access Manager. Enter this option as <i>yes</i> only when the Tivoli Access Manager domain is in an irreparable state.
Defer	Enter <i>no</i> , to force the unconfiguration of the connected server. Enter <i>No</i> for the unconfiguration to proceed correctly.

4. When all information is entered, enter F to save the properties or C to cancel from the unconfiguration process and discard the entered information.
5. Optional: Synchronize all nodes.
6. Restart all WebSphere Application Server instances for the changes to take effect.

Configuring security auditing using scripting

Security auditing provides tracking and archiving of auditable events. This topic uses the wsadmin tool to enable and administer your security auditing configurations.

About this task

While security authentication and authorization ensures that users must have access to view protected resources, security auditing provides a mechanism to validate the integrity of a security computing environment. Security auditing collects and logs authentication, authorization, system management, security, and audit policy events in audit event records. You can analyze audit event records to determine possible security breaches, threats, attacks, and potential weaknesses in the security configuration of your environment. Enable security auditing in your environment. For example, the following list displays a sample of events to audit:

- Determine the time that a specific user attempted to access a resource.
- View information for successful and unsuccessful attempts to access resources.
- Review changes to resources that were made by a specific user.
- Determine the cause of unsuccessful login attempts.

Use the following task outline to enable and configure security auditing in your environment:

1. Enable administrative security in your environment.
2. Configure auditable events. The security auditing configuration provides four default auditable filters. Use this topic to configure filters for additional audit events.
3. Configure audit event factories. The security auditing configuration provides a default event factory. Use this topic to configure additional audit event factories.
4. Configure audit service providers. The security auditing configuration provides a default service provider. Use this topic to configure additional audit service providers.
5. Set the global audit policy. After setting up audit event factories, service providers, and events, use this topic to enable security auditing.

Results

After completing the steps to enable and configure security auditing, the profile of interest audits your security configurations for specific auditable event types.

What to do next

To further configure security auditing, you can:

- Configure audit notifications.
- Encrypt audit records.
- Sign audit records.

Configuring audit service providers using scripting

Before enabling security auditing, use this task to configure audit service providers using the wsadmin tool. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before configuring security audit service providers, enable administrative security in your environment.

About this task

In order to enable security auditing in your environment, you must configure an audit service provider. The audit service provider writes the audit records and data to the back-end repository associated with the service provide implementation. The security auditing configuration provides a default service provider. Use this topic to customize your security auditing subsystem by creating additional audit service providers.

Use the following steps to configure your security auditing subsystem using the wsadmin tool:

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Configure an audit service provider. You can use the default binary-based audit service provider, or use this step to create a new audit service provider.

There are binary file-based and third-party audit service providers. In addition to the default binary file-based service provider, you can configure a third-party audit service provider.

Choose the type of audit service provider to create.

- Use the createBinaryEmitter command and the following required parameters to create a default audit service provider:

Parameter	Description	Data Type	Required
-uniqueName	Specifies a unique name that identifies the audit service provider.	String	Yes
-className	Specifies the class implementation of the audit service provider interface.	String	Yes
-fileLocation	Specifies the file location for the audit service provider to write the audit logs.	String	Yes
-auditFilters	Specifies a reference or a group of references to predefined audit filters, using the following format: reference, reference, reference	String	Yes
-maxFileSize	Specifies the maximum size each audit log reaches before the system saves it with a timestamp and creates a new file. Specify the file size in megabytes. If you do not specify this parameter, the system sets the maximum file size to 10 megabytes.	Integer	No
-maxLogs	Specifies the maximum number of audit logs to create before rewriting the oldest log. If you do not specify this parameter, the system allows up to 100 audit logs before overwriting the oldest log.	Integer	No

The following example creates a new audit service provider in your security auditing configuration:

```
AdminTask.createBinaryEmitter('-uniqueName newASP -className  
com.ibm.ws.security.audit.BinaryEmitterImpl -fileLocation /AUDIT_logs  
-auditFilters "AuditSpecification_1173199825608, AuditSpecification_1173199825609,  
AuditSpecification_1173199825610, AuditSpecification_1173199825611"')
```


- Use the `createThirdPartyEmitter` command to use a third-party audit service provider. Use the following parameters with the `createThirdPartyEmitter` command:

Parameter	Description	Data Type	Required
-uniqueName	Specifies a unique name that identifies the audit service provider.	String	Yes
-className	Specifies the class implementation of the audit service provider interface.	String	Yes
-eventFormatterClass	Specifies the class that implements how the audit event is formatted for output. If you do not specify this parameter, the system uses the standard text format for output.	String	Yes
-auditFilters	Specifies a reference identifier or a group of reference identifiers to pre-defined audit filters, using the following format: reference, reference, reference.	String	Yes
-customProperties	Specifies any custom properties that might be required to configure a third party audit service provider.	String	No

The following example creates a new third party audit service provider in your security auditing configuration:

```
AdminTask.createThirdPartyEmitter('-uniqueName myAuditServiceProvider -className
com.mycompany.myclass -fileLocation /auditLogs -auditFilters
"AuditSpecification_1173199825608, AuditSpecification_1173199825609,
AuditSpecification_1173199825610, AuditSpecification_1173199825611"')
```

3. Save your configuration changes.

What to do next

Enable security auditing in your environment.

Configuring audit event factories using scripting

Before enabling security auditing, use this task to configure audit event factories using the `wsadmin` tool. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before configuring security auditing event factories, enable administrative security in your environment.

About this task

In order to enable security auditing in your environment, you must configure an audit event factory. The audit event factory gathers the data that is associated with security events. The security auditing configuration provides a default event factory. Use this topic to customize your security auditing subsystem by creating additional audit event factories.

Use the following steps to configure your security auditing subsystem using the `wsadmin` tool:

1. Launch the `wsadmin` scripting tool using the Jython scripting language.
2. Configure event filters. You can use the default event filters or use this step to create additional filters to customize your security auditing configuration.

The application server provides the following event filters by default in the `audit.xml` template file:

Event Name	Outcome of event
SECURITY_AUTHN	SUCCESS
SECURITY_AUTHN	DENIED
SECURITY_RESOURCE_ACCESS	SUCCESS
SECURITY_AUTHN	REDIRECT

You can configure additional audit event types to track and archive various events. Use the following command to list all supported auditable events:

```
print AdminTask.getSupportedAuditEvents()
```

Use the `createAuditFilter` command with the `-eventType` and `-outcome` parameters to enable one or multiple audit events and outcomes. You can specify multiple event types and multiple outcomes separated by a comma with one command invocation. The following list describes each valid auditable event that you can specify with the `-eventType` parameter:

Event name	Description
SECURITY_AUTHN	Audits all authentication events
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities are involved
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation of an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might be related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that might represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web services
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including identity assertion, RunAs, and low assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity

For each audit event type, you must specify an outcome. Valid outcomes include **SUCCESS**, **FAILURE**, **REDIRECT**, **ERROR**, **DENIED**, **WARNING**, and **INFO**. The following command example creates an audit filter to log users who receive an error when modifying credentials:

```
AdminTask.createAuditFilter('-name uniqueFilterName -eventType
SECURITY_AUTHN_CREDS_MODIFY,SECURITY_AUTHN_DELEGATION -outcome ERROR,REDIRECT')
```

3. Create an audit event factory. You can use the default audit event factory or use this step to create a new audit event factory.

Use the `createAuditEventFactory` command to create an audit event factory in your security configuration. You can use the default implementation of the audit event factory or use a third-party implementation. To configure a third-party implementation, use the optional `-customProperties` parameter to specify any properties necessary to configure the audit event factory implementation.

Specify the following required parameters with the `createAuditEventFactory` to configure your audit event factory:

Parameter	Description	Data type	Required
-uniqueName	Specifies a unique name that identifies the audit event factory.	String	Yes
-className	Specifies the class implementation of the audit event factory interface.	String	Yes

Parameter	Description	Data type	Required
-auditFilters	Specifies a reference or a group of references to predefined audit filters, using the following format: "reference, reference, reference"	String	Yes
-provider	Specifies a reference to a predefined audit service provider implementation.	String	Yes
-customProperties	Specifies a comma (,) separated list of custom property pairs to add to the security object in the following format: attribute=value,attribute=value	String	No

The following sample command creates an enables an audit event factory:

```
AdminTask.createAuditEventFactory('-uniqueName eventFactory1 -className
com.ibm.ws.security.audit.AuditEventFactoryImpl -auditFilters
"AuditSpecification_1173199825608, AuditSpecification_1173199825609, AuditSpecification_1173199825610,
AuditSpecification_1173199825611" -provider newASP')
```

4. Save your configuration changes.

What to do next

Configure the audit service provider.

Configuring auditable events using scripting

Before enabling security auditing, use this task to configure event filters using the wsadmin tool. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before configuring security auditing filters, enable administrative security in your environment.

About this task

Before configuring an audit event factory and audit service provider, configure event filters. The audit service provider writes audit records to the back end repository associated with the provider implementation. The audit event factory generates security events. Event filters specify which event types and outcomes the system audits and records. Each event type has up to seven possible outcomes, including success, failure, denied, error, warning, info, and redirect. The security auditing configuration provides four default filters. Use this topic to customize your security auditing subsystem by creating additional audit event filters.

Use the following steps to configure your security auditing subsystem using the wsadmin tool:

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Configure event filters. You can use the default event filters or use this step to create additional filters to customize your security auditing configuration.

The application server provides the following event filters by default in the `audit.xml` template file:

Event Name	Outcome of event
SECURITY_AUTHN	SUCCESS
SECURITY_AUTHN	DENIED
SECURITY_RESOURCE_ACCESS	SUCCESS
SECURITY_AUTHN	REDIRECT

You can configure additional audit event types to capture various events. Use the following command to list all supported auditable events:

```
print AdminTask.getSupportedAuditEvents()
```

Use the `createAuditFilter` command with the `-name`, `-eventType`, and `-outcome` parameters to enable one or multiple audit events and outcomes. You can specify multiple event types and multiple

outcomes separated by a comma with one command invocation. The following list describes each valid auditable event that you can specify with the `-eventType` parameter:

Event name	Description
SECURITY_AUTHN	Audits all authentication events
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities are involved
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including identity assertion, RunAs, and low assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.

For each audit event type, you must specify an outcome. Valid outcomes include `SUCCESS`, `FAILURE`, `REDIRECT`, `ERROR`, `DENIED`, `WARNING`, and `INFO`. The following command example creates an audit filter to log users who receive an error when modifying credentials:

```
AdminTask.createAuditFilter('-name myUniqueName -eventType SECURITY_AUTHN_CREDS_MODIFY,SECURITY_AUTHN_DELEGATION -outcome ERROR,REDIRECT')
```

3. Save your configuration changes.

What to do next

Enable security auditing in your environment.

Enabling security auditing using scripting

Use this task to enable and configure security auditing in your environment with the `wsadmin` tool. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before enabling security auditing, enable administrative security in your environment.

If you previously configured security auditing and do not want to modify configuration settings, use the `enableAudit` and `disableAudit` commands to start and stop security auditing. After enabling or disabling security auditing, restart the server to apply the configuration changes.

About this task

Security auditing ensures the integrity of a security computing environment. Security auditing collects and logs authentication, authorization, system management, security, and audit policy events in audit event records. You can analyze audit event records to determine possible security breaches, threats, attacks, and potential weaknesses in the security configuration of your environment.

Use the following steps to enable and configure security auditing in your environment:

1. Launch the `wsadmin` scripting tool using the Jython scripting language.
2. Verify that the security auditing subsystem is configured.

To enable security auditing, you must configure event filters, an audit emitter, and an audit event factory. Event filters specify which event types the system audits and records, and the outcome of the event. The audit service provider writes the audit records to the backend repository that is associated with the implementation. The audit event factory generates security events.

By default, the security auditing system includes one audit service provider and one audit event factory.

The audit command groups provide several commands to query for event filters, audit emitters, event factories, and their respective configuration attributes. Use the audit command reference to use specific query commands. The following example commands query your security auditing configuration at a high level.

- Use the `getAuditFilters` command to display a list of references to all audit filters defined in your configuration, as the following example demonstrates:

```
AdminTask.getAuditFilters()
```

- Use the `listAuditEmitters` command to display a list of all audit emitters in your configuration, as the following example demonstrates:

```
AdminTask.listAuditEmitters()
```

- Use the `listAuditEventFactories` command to display a list of all audit event factories in your configuration, as the following example demonstrates:

```
AdminTask.listAuditEventFactories()
```

3. Enable security auditing in your environment. Use the `modifyAuditPolicy` command to enable security auditing in your environment. Use the following optional parameters for the `modifyAuditPolicy` command to customize your security auditing configuration:

Parameter	Description	Data type	Required
<code>-auditEnabled</code>	Specifies whether to enable security auditing.	Boolean	No
<code>-auditPolicy</code>	Specifies the behavior of the server process if the audit subsystem fails. Valid values are: <code>WARN</code> , <code>NOWARN</code> and <code>FATAL</code> . The <code>WARN</code> setting notifies the auditor when an error occurs and ceases auditing when an error occurs in the audit sub-system, but continues to run the application server process. The <code>NOWARN</code> setting does not notify the auditor when an error occurs and ceases auditing, but continues to run the application server process. The <code>FATAL</code> setting notifies the auditor of the error and stops the application server process. By default, the command assigns the <code>NOWARN</code> setting.	String	No
<code>-auditorId</code>	Specifies the ID of the user to assign to the auditor role.	String	No
<code>-auditorPwd</code>	Specifies the password for the auditor role.	String	No
<code>-sign</code>	Specifies whether to sign audit records. By default, the security auditing system does not sign audit records. You must configure the signing of audit records before you can specify this parameter.	Boolean	No
<code>-encrypt</code>	Specifies whether to encrypt audit records. By default, the security auditing system does not encrypt audit records. You must configure encryption for audit records before you can specify this parameter.	Boolean	No
<code>-verbose</code>	Specifies whether to capture verbose audit data. By default, the security auditing system does not capture verbose audit data.	Boolean	No
<code>-encryptionCert</code>	Specifies the reference ID of the certificate to use for encryption. Specify this parameter if you set the <code>-encrypt</code> parameter to <code>true</code> .	String	No

The following example command enables security auditing, and identifies the primary auditor by assigning a user and password.

```
AdminTask.modifyAuditPolicy('-auditEnabled true -auditorId securityAdmin -auditorPwd security4you')
```

4. Save your configuration changes.
5. Restart the server.

Results

After completing the steps to enable and configure security auditing, the profile of interest audits your security configurations for specific auditable event types.

What to do next

After you configure the audit policy for the first time, use the `enableAudit` and `disableAudit` commands to turn the security auditing system on and off. The system maintains the settings that you define with the `modifyAuditPolicy` command when you enable and disable the security auditing system.

Note: You must restart the server to apply the configuration changes.

Configuring security audit notifications using scripting

Configure the security auditing system to send email notifications to a distribution list, system log, or both a distribution list and a system log if a failure occurs in the audit subsystem. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before configuring a notification object in the `audit.xml` configuration file, verify that you set up a security auditing subsystem and configured the security auditing policy.

About this task

You can configure the security auditing system to notify a specific person or group when a failure occurs in the audit subsystem. Use the following steps to enable security auditing email notifications, set the format of notification email, and secure email:

1. Launch the `wsadmin` scripting tool using the Jython scripting language.
2. Customize and enable security auditing email notifications.

Use the `createAuditNotification` command and the following parameters to configure notifications:

Parameter	Description	Data Types	Required
<code>-notificationName</code>	Specifies a unique name to assign the audit notification object in the <code>audit.xml</code> file.	String	Yes
<code>-logToSystemOut</code>	Specifies whether to log the notification to the <code>SystemOut.log</code> file.	Boolean	Yes
<code>-sendEmail</code>	Specifies whether to email notifications.	Boolean	Yes
<code>-emailList</code>	Specifies the email address or email distribution list to email notifications. The format for this parameter is: <code>admin@company.com(smtp-server.mycompany.com)</code>	String	No
<code>-emailFormat</code>	Specifies whether to send the email be HTML or TEXT format.	String	No

To create the audit notification object, you must specify the `-notificationName`, `-logToSystemOut`, and `-sendEmail` parameters, as the following example demonstrates:

```
AdminTask.createAuditNotification('-notificationName defaultEmailNotification
-logToSystemOut true -sendEmail true -emailList administrator@mycompany.com(smtp-server.mycompany.com)
-emailFormat HTML')
```

3. Create an audit notification monitor object.

Create an audit notification monitor object to monitor the security auditing subsystem for possible failure. Use the `createAuditNotificationMonitor` command and the following parameters to create a monitor object for the security auditing system:

Parameter	Description	Data Types	Required
-notificationName	Specifies a unique name to assign the audit notification object in the audit.xml file.	String	Yes
-logToSystemOut	Specifies whether to log the notification to the SystemOut.log file.	Boolean	Yes
-sendEmail	Specifies whether to email notifications.	Boolean	Yes
-emailList	Specifies the email address or email distribution list to email notifications. The format for this parameter is: admin@company.com(smtp-server.mycompany.com)	String	No
-emailFormat	Specifies whether to send the email be HTML or TEXT format.	String	No

To create the audit notification monitor object, you must specify the `-notificationName`, `-logToSystemOut`, and `-sendEmail` parameters, as the following example demonstrates:

```
AdminTask.createAuditNotificationMonitor('-notificationName defaultEmailNotification  
-logToSystemOut true -sendEmail true -emailList administrator@mycompany.com(smtp-server.mycompany.com)  
-emailFormat HTML')
```

4. Save your configuration changes.

Results

The security auditing system notifies the specified recipients if a failure occurs in the security auditing system.

What to do next

Use the `modifyAuditNotification` command and the Audit Notification Commands command group for the `AdminTask` object to manage your notification configuration.

Encrypting security audit data using scripting

You can use the `wsadmin` tool to configure the security auditing system to encrypt security audit records. Security auditing provides tracking and archiving of auditable events.

Before you begin

Before configuring encryption, set up your security auditing subsystem. You can enable security auditing before or after completing the steps in this topic.

Verify that you have the appropriate administrative role. To complete this topic, you must have the auditor administrative role. If you are importing a certificate from a keystore that exists in the `security.xml` file, you must have the auditor and administrator administrative roles.

About this task

When configuring encryption, the auditor can select one of the following choices:

- Allow the application server to automatically generate a certificate or use an existing self-signed certificate generated by the auditor.

- Use an existing keystore to store this certificate, or create a new keystore to store this certificate.

Note: To ensure that there is a separation of privileges between the administrator role and the auditor role, the auditor can create a self-signed certificate outside of the application server process and maintain the private key of that certificate.

Use the following task steps to encrypt security audit data:

1. Launch the wsadmin scripting tool using the Jython scripting language.
2. Configure encryption settings for security audit data.

Use the createAuditEncryptionConfig command and the following parameters to create the audit encryption model to encrypt your audit records. You must specify the -enableAuditEncryption, -certAlias, and -encryptionKeyStoreRef parameters, and either the -autogenCert or -importCert parameters.

Parameter	Description	Data Type	Required
-enableAuditEncryption	Specifies whether to encrypt audit records. This parameter modifies your audit policy configuration.	Boolean	Yes
-certAlias	Specifies the alias name that identifies the generated or imported certificate.	String	Yes
-encryptionKeyStoreRef	Specifies the reference ID of the keystore to import the certificate to.	String	Yes
-autogenCert	Specifies whether to automatically generate the certificate used to encrypt the audit records. You must specify either this parameter or the -importCert parameter, but you cannot specify both.	Boolean	No
-importCert	Specifies whether to import an existing certificate to encrypt the audit records. You must specify either this parameter or the -autogenCert parameter, but you cannot specify both.	Boolean	No
-certKeyFileName	Specifies the unique name of the key file from which the certificate is imported.	String	No
-certKeyFilePath	Specifies the key file location from which the certificate is imported.	String	No
-certKeyFileType	Specifies the key file type from which the certificate is imported.	String	No
-certKeyFilePassword	Specifies the key file password from which the certificate is imported.	String	No
-certAliasToImport	Specifies the alias from which the certificate is imported.	String	No

The following command example configures encryption and supports the system to automatically generate the certificate:

```
AdminTask.createAuditEncryptionConfig('-enableAuditEncryption true -certAlias auditCertificate
-autogenCert true -encryptionKeyStoreRef auditKeyStore')
```

The following command example configures encryption and imports a certificate:

```
AdminTask.createAuditEncryptionConfig('-enableAuditEncryption true -certAlias auditCertificate
-importCert true -certKeyFileName MyServerKeyFile.p12 -certKeyFilePath
install_root/etc/MyServerKeyFile.p12 -certKeyFileType PKCS12 -certKeyFilePassword password4key
-certAliasToImport defaultCertificate -encryptionKeyStoreRef auditKeyStore')
```

3. You must restart the server to apply configuration changes.

Results

Encryption is configured for security audit data. If you set the -enableAuditEncryption parameter to true, then your security auditing system encrypts security audit data when security auditing is enabled.

What to do next

After you configure the encryption model for the first time, then you may use the `enableAuditEncryption` and `disableAuditEncryption` commands to turn encryption on and off.

The following example uses the `enableAuditEncryption` command to turn on encryption:

```
AdminTask.enableAuditEncryption()
```

The following example uses the `disableAuditEncryption` command to turn off encryption:

```
AdminTask.disableAuditEncryption()
```

Signing security audit data using scripting

You can use the `wsadmin` tool to configure the security auditing system to sign security audit records. Security auditing provides tracking and archiving of auditable events.

Before you begin

Verify that you have the appropriate administrative role. To complete this topic, you must have the auditor and administrator administrative roles.

About this task

When configuring the signing of audit data, the auditor can choose between the following options:

- Allow the application server to automatically generate a certificate.
- Use an existing self-signed certificate that the auditor previously generated.
- Use the same self-signed certificate as the system uses to encrypt the audit records.
- Use an existing keystore to store this certificate.
- Create a new keystore to store this certificate.
- Use an existing self-signed certificate in an existing keystore.

Use the following task steps to configure the signing of security audit data:

1. Launch the `wsadmin` scripting tool using the Jython scripting language.
2. Configure signing settings for security audit data.

Use the `createAuditSigningConfig` command to create the signing model to sign your audit records.

You can import the certificate from an existing key file name that contains that certificate, automatically generate the certificate, or use the same certificate as used to encrypt the audit records. The signing keystore must exist in the `security.xml` file. The system updates this keystore with the certificate to use to sign the audit records. Use the parameters in the following table with the `createAuditSigningConfig` command. You must specify the `-enableAuditSigning`, `-certAlias`, and `-signingKeyStoreRef` parameters.

Parameter	Description	Data Type	Required
<code>-enableAuditSigning</code>	Specifies whether to sign audit records. This parameter modifies your audit policy configuration.	Boolean	Yes
<code>-certAlias</code>	Specifies the alias name that identifies the generated or imported certificate.	String	Yes
<code>-signingKeyStoreRef</code>	Specifies the reference ID of the keystore to import the certificate to.	String	Yes
<code>-useEncryptionCert</code>	Specifies whether to use the same certificate for encryption and signing. You must specify the <code>-useEncryptionCert</code> , <code>-autogenCert</code> , or <code>-importCert</code> parameter.	Boolean	No

Parameter	Description	Data Type	Required
-autogenCert	Specifies whether to automatically generate the certificate used to sign the audit records. You must specify the -useEncryptionCert, -autogenCert, or -importCert parameter.	Boolean	No
-importCert	Specifies whether to import an existing certificate to sign the audit records. You must specify the -useEncryptionCert, -autogenCert, or -importCert parameter.	Boolean	No
-certKeyFileName	Specifies the unique name of the key file for the certificate to import.	String	No
-certKeyFilePath	Specifies the key file location for the certificate to import.	String	No
-certKeyFileType	Specifies the key file type for the certificate to import.	String	No
-certKeyFilePassword	Specifies the key file password for the certificate to import.	String	No
-certAliasToImport	Specifies the alias of the certificate to import.	String	No

The following command example configures signing and allows the system to automatically generate the certificate:

```
AdminTask.createAuditSigningConfig('-enableAuditSigning true -certAlias auditSigningCert
-autogenCert true -signingKeyStoreRef Ref_Id_of_KeyStoreInSecurityXML')
```

The following command example configures signing and imports a certificate:

```
AdminTask.createAuditSigningConfig('-enableAuditSigning true -certAlias auditSigningCert
-importCert true -certKeyFileName MyServerKeyFile.p12 -certKeyFilePath install_root/etc/MyServerKeyFile.p12
-certKeyFileType PKCS12 -certKeyFilePassword password4key -certAliasToImport defaultCertificate
-signingKeyStoreRef Ref_Id_of_KeyStoreInSecurityXML')
```

The following command example uses the same certificate for signing and encryption:

```
AdminTask.createAuditSigningConfig('-enableAuditSigning true -certAlias auditSigningCert
-useEncryptionCert true -signingKeyStoreRef Ref_Id_of_KeyStoreInSecurityXML')
```

3. Save your configuration changes.
4. Restart the server to apply the configuration changes.

Results

Signing is configured for your security audit data. If you set the -enableAuditSigning parameter to true, your security auditing system signs security audit data when security auditing is enabled.

What to do next

Once you configure the signing model for the first time, use the enableAuditSigning and disableAuditSigning commands to quickly turn signing on and off. The following example uses the enableAuditSigning command to turn signing on:

```
AdminTask.enableAuditSigning()
```

The following example uses the disableAuditSigning command to turn signing off:

```
AdminTask.disableAuditSigning()
```

AuditKeyStoreCommands command group for the AdminTask object

You can use the Jython scripting language to configure the security auditing system with the wsadmin tool. Use the commands and parameters in the AuditKeyStoreCommands group to configure audit keystores in the security auditing system.

Use the following commands to manage audit key stores in the audit.xml configuration file:

- createAuditKeyStore

- deleteAuditKeyStore
- getAuditKeyStoreInfo
- listAuditKeyStores
- modifyAuditKeyStore

createAuditKeyStore

Creates a keystore in the `audit.xml` file. The system uses this keystore to encrypt audit records.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

-keyStoreType

Specifies a valid keystore type. The default keystore type is PKCS12. (String, required)

-keyStoreLocation

Specifies the location where the system creates the keystore. (String, required)

-keyStorePassword

Specifies the password for the keystore. (String, required)

-keyStorePasswordVerify

Verifies the password for the keystore. (String, required)

Optional parameters

-keyStoreProvider

Specifies a provider for the keystore. (String, optional)

-keyStoreIsFileBased

Specifies if the keystore is file-based. The default is `true`. (Boolean, optional)

-keyStoreHostList

Specifies the host list for the keystore. (String, optional)

-keyStoreInitAtStartup

Specifies whether the system initializes the keystore on startup. The default is `false`. (Boolean, optional)

-keyStoreReadOnly

Specifies whether the keystore is read-only or not. Default is `false`. (Boolean, optional)

-keyStoreStashFile

Specifies whether the keystore needs a stash file. (Boolean, optional)

-enableCryptoOperations

Specifies whether the keystore is an acceleration keystore. False default. (Boolean, optional)

-scopeName

Specifies the scope for the keystore. (String, optional)

-keyStoreDescription

Specifies a description for the keystore. (String, optional)

Return value

The command returns the ID of the new keystore, as the following example displays:

```
KeyStore_1173199825578
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditKeyStore('-keyStoreName mynewkeystore -keyStoreLocation  
c:\install_root\appserver\profiles\AppSrv01\config\cells -keyStorePassword  
myPwd -keyStorePasswordVerify myPwd -keyStoreProvider IBMJCE -scopeName (cell):Node04Cell')
```

- Using Jython list:

```
AdminTask.createAuditKeyStore(['-keyStoreName', 'mynewkeystore', '-keyStoreLocation',  
'c:\install_root\appserver\profiles\AppSrv01\config\cells', '-keyStorePassword',  
'myPwd', '-keyStorePasswordVerify', 'myPwd', '-keyStoreProvider', 'IBMJCE',  
'-scopeName', '(cell):Node04Cell'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditKeyStore('-interactive')
```

deleteAuditKeyStore

The `deleteAuditKeyStore` command removes the reference to an audit keystore from the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the name of the keystore. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the keystore. (String, optional)

-removeKeyStoreFile

Specifies whether to remove the keystore from the configuration. Specify this parameter if the keystore of interest is not in use. (Boolean, optional)

Return value

The command returns a value of `true` if the system successfully removes the reference to the keystore from the `audit.xml` configuration file.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditKeyStore('-keyStoreName AuditDefaultKeyStore -scopeName  
(cell):Node04Cell -removeKeyStoreFile false')
```

- Using Jython list:

```
AdminTask.deleteAuditKeyStore(['-keyStoreName', 'AuditDefaultKeyStore', '-scopeName',  
'(cell):Node04Cell', '-removeKeyStoreFile', 'false'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditKeyStore('-interactive')
```

getAuditKeyStoreInfo

The `getAuditKeyStoreInfo` command returns a list of attributes for the keystore that the system uses to encrypt audit records.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name to identify the keystore. (String, required)

Optional parameters

-scopeName

Specifies the management scope of the keystore. (String, optional)

Return value

The command returns a list of attributes for the keystore, as the following sample output displays:

```
{location ${CONFIG_ROOT}/audittrust.p12}
{password *****}
{websphere_config_data_id cells/Node04Cell|audit.xml#KeyStore_1173199825578}
{websphere_config_data_version {}}
{useforacceleration false}
{slot 0}
{type PKCS12}
{additionalkeystoreattrs {}}
{filebased true}
{websphere_config_data_type KeyStore}
{customproviderclass {}}
{hostlist {}}
{createstashfileforcms false}
{description {keyStore description}}
{readonly false}
{initializeatstartup true}
{managementscope (cells/Node04Cell|audit.xml#ManagementScope_1173199825608)}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditKeyStoreInfo('-keyStoreName AuditDefaultKeyStore')
```

- Using Jython list:

```
AdminTask.getAuditKeyStoreInfo(['-keyStoreName', 'AuditDefaultKeyStore'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditKeyStoreInfo('-interactive')
```

listAuditKeyStores

The `listAuditKeyStores` command lists the attributes for the audit keystores within a specific management scope or for all audit keystores.

The user must have the monitor administrative role to run this command.

Target object

None.

Optional parameters

-scopeName

Specifies the management scope associated with the keystores of interest. (String, optional)

-all

Specifies whether to list all keystores. When the -all parameter is set as true, it overrides the -scopeName parameter. (Boolean, optional)

Return value

The command returns a list of attributes for the scope of interest, as the following sample output displays:

```
{{location ${CONFIG_ROOT}/audittrust.p12}
{password *****}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#KeyStore_1173199825578}
{_Websphere_Config_Data_Version {}}
{useForAcceleration false}
{slot 0}
{type PKCS12}
{additionalKeyStoreAttrs {}}
{fileBased true}
{_Websphere_Config_Data_Type KeyStore}
{customProviderClass {}}
{hostList {}}
{keyStoreRef KeyStore_1173199825578}
{createStashFileForCMS false}
{description {keyStore description}}
{managementScope (cells/Node04Cell|audit.xml#ManagementScope_1173199825608)}
{readOnly false}
{initializeAtStartup true}
{usage {}}
{provider IBMJCE}{name AuditDefaultKeyStore}}
{{location c:\install_root\appserver\profiles\AppSrv01\config\cells}
{password *****}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#KeyStore_1184700968484}
{_Websphere_Config_Data_Version {}}
{useForAcceleration false}
{slot 0}
{type PKCS12}
{additionalKeyStoreAttrs {}}
{fileBased true}
{_Websphere_Config_Data_Type KeyStore}
{customProviderClass {}}
{hostList {}}
{keyStoreRef KeyStore_1184700968484}
{createStashFileForCMS false}
{description {}}
{managementScope {}}
{readOnly false}
{initializeAtStartup false}
{usage {}}
{provider IBMJCE}
{name mykeystore}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditKeyStores('-scopeName (cell):Node04Cell')
```

- Using Jython list:

```
AdminTask.listAuditKeyStores(['-scopeName', '(cell):Node04Cell'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditKeyStores('-interactive')
```


modifyAuditKeyStore

The `modifyAuditKeyStore` command modifies the keystore reference in the `audit.xml` file. The command edits keystore that encrypts audit records.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

Optional parameters

-scopeName

Specifies the scope name of this keystore. (String, optional)

-keyStoreType

Specifies valid keystore type. (String, optional)

-keyStoreLocation

Specifies the location where the system creates the keystore. (String, optional)

-keyStorePassword

Specifies the password for this keystore. (String, optional)

-keyStoreIsFileBased

Specifies whether the keystore is file based. (Boolean, optional)

-keyStoreInitAtStartup

Specifies whether the system should initialize the keystore at startup. (Boolean, optional)

-keyStoreReadOnly

Specifies whether the keystore is read-only or editable. (Boolean, optional)

-keyStoreDescription

Specifies a description for the keystore. (String, optional)

Return value

The command returns a value of `true` if the system successfully modifies the keystore.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditKeyStore('-keyStoreName AuditDefaultKeyStore -scopeName
(cell):Node04Cell -keyStoreType PKCS12 -keyStoreLocation
c:\install_root\appserver\profiles\AppSrv01\config\cells\Node04Cell\audittrust.p12
-keyStorePassword myPwd')
```

- Using Jython list:

```
AdminTask.modifyAuditKeyStore(['-keyStoreName', 'AuditDefaultKeyStore', '-scopeName',
'(cell):Node04Cell', '-keyStoreType', 'PKCS12', '-keyStoreLocation',
'c:\install_root\appserver\profiles\AppSrv01\config\cells\Node04Cell\audittrust.p12',
'-keyStorePassword', 'myPwd'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditKeyStore('-interactive')
```

AuditEmitterCommands for the AdminTask object

You can use the Jython scripting language to configure audit service providers with the wsadmin tool. Use the commands and parameters in the AuditEmitterCommands group to create, manage, and remove audit service providers from your security auditing system configuration.

Use the following commands to configure audit service providers:

- “createBinaryEmitter”
- “createSMFEmitter” on page 1078
- “createThirdPartyEmitter” on page 1079
- “deleteAuditEmitterByRef” on page 1080
- “deleteAuditEmitterByName” on page 1080
- “getAuditEmitter” on page 1081
- “getBinaryFileLocation” on page 1082
- “getAuditEmitterFilters” on page 1082
- “getBinaryFileSize” on page 1083
- “getEmitterClass” on page 1083
- “getEmitterUniqueld” on page 1084
- “getMaxNumBinaryLogs” on page 1084
- “listAuditEmitters” on page 1085
- “modifyAuditEmitter” on page 1086
- “setAuditEmitterFilters” on page 1087

createBinaryEmitter

The createBinaryEmitter command creates an entry in the `audit.xml` file to reference the configuration of the binary file emitter implementation of the audit service provider interface.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies a name to uniquely identify this implementation of the audit service provider interface. (String, required)

-className

Specifies the class that implements the audit service provider interface. (String, required)

-fileLocation

Specifies the location where the system writes the audit logs. (String, required)

-auditFilters

Specifies a reference or a group of references to predefined audit filters. Use the following format to specify multiple references: `reference,reference,reference` (String, required)

Optional parameters

-eventFormatterClass

Specifies the class that implements how the system formats the audit event for output. If you want to use the default audit service provider, do not specify this parameter. (String, optional)

-maxFileSize

Specifies the maximum size that each log reaches before the system saves the audit log with a timestamp. Specify the size in megabytes. The default value is 10 MB. (Integer, optional)

-maxLogs

Specifies the maximum number of log files to create before the system rewrites the oldest audit log. The default value is 100 logs. (Integer, optional)

Return value

The command returns the shortened reference ID for the audit service provider, as the following sample output displays:

```
AuditServiceProvider_1184686384968
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createBinaryEmitter('-uniqueName mybinaryemitter -className  
com.ibm.ws.security.audit.BinaryEmitterImpl -fileLocation  
c:\wasinstall\appserver\profiles\AppSrv01\logs\server1 -maxFileSize 20 -maxLogs  
100 -auditFilters AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.createBinaryEmitter(['-uniqueName', 'mybinaryemitter', '-className',  
'com.ibm.ws.security.audit.BinaryEmitterImpl', '-fileLocation',  
'c:\wasinstall\appserver\profiles\AppSrv01\logs\server1', '-maxFileSize',  
'20', '-maxLogs', '100', '-auditFilters', 'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createBinaryEmitter('-interactive')
```

createSMFEmitter

The createSMFEmitter command creates an entry in the audit.xml file to reference the configuration of an SMF implementation of the audit service provider interface. The encryption and signing of audit records is not supported for SMF implementations.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies a name to uniquely identify this implementation of the audit service provider interface. (String, required)

-auditFilters

Specifies a reference or a group of references to predefined audit filters. Use the following format to specify multiple references: reference,reference,reference (String, required)

Return value

The command returns the shortened reference ID for the audit service provider, as the following sample output displays:

```
AuditServiceProvider_1184686384968
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createSMFEmitter('-uniqueName mySMFEmitter -auditFilters
AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.createSMFEmitter(['-uniqueName', 'mySMFEmitter', '-auditFilters',
'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createSMFEmitter('-interactive')
```

createThirdPartyEmitter

The `createThirdPartyEmitter` command creates an entry in the `audit.xml` configuration file to reference the configuration of a third party emitter implementation of the audit service provider interface. The encryption and signing of audit records is not supported for third party implementations.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies a name to uniquely identify this implementation of the audit service provider interface. (String, required)

-className

Specifies the class that implements the audit service provider interface. (String, required)

-auditFilters

Specifies a reference or a group of references to predefined audit filters. Use the following format to specify multiple references: `reference,reference,reference` (String, required)

Optional parameters

-eventFormatterClass

Specifies the class that implements how the system formats the audit event for output. (String, optional)

-customProperties

Specifies any custom properties that the system might need to configure the third party implementation of the audit service provider. Use the following format to specify the custom properties: `name=value,name=value` (String, optional)

Return value

The command returns the shortened reference ID to the audit service provider, as the following example output displays:

```
AuditServiceProvider_1184686638218
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createThirdPartyEmitter('-uniqueName myThirdPartyEmitter -className
com.mycompany.myemitterclass -eventFormatterClass com.mycompany.myeventformatterclass
-auditFilters AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.createThirdPartyEmitter(['-uniqueName', 'myThirdPartyEmitter', '-className',  
'com.mycompany.myemitterclass', '-eventFormatterClass', 'com.mycompany.myeventformatterclass',  
'-auditFilters', 'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createThirdPartyEmitter('-interactive')
```

deleteAuditEmitterByRef

The `deleteAuditEmitterByRef` command deletes the audit service provider implementation that the system references with the reference id. If an event factory is using the audit service provider, the system generates an error that indicates that the system cannot remove the audit service provider.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies the reference identifier of the audit service provider implementation to delete. (String, required)

Return value

The command returns a value of `true` if the system successfully removes the audit service provider.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditEmitterByRef('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.deleteAuditEmitterByRef(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditEmitterByRef('-interactive')
```

deleteAuditEmitterByName

The `deleteAuditEmitterByName` command deletes the audit service provider implementation that the system references with the unique name. If an event factory is using the audit service provider, the system generates an error that indicates that the system cannot remove the audit service provider.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies the name that uniquely identifies this implementation of the audit service provider interface to delete. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit service provider implementation.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditEmitterByName('-uniqueName mybinaryemitter')
```

- Using Jython list:

```
AdminTask.deleteAuditEmitterByName(['-uniqueName', 'mybinaryemitter'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditEmitterByName('-interactive')
```

getAuditEmitter

The `getAuditEmitter` command returns the attributes for the audit service provider of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a audit service provider implementation. (String, required)

Return value

The command returns an attribute list for the audit service provider specified by the `-emitterRef` parameter, as the following example output displays:

```
{auditSpecifications myfilter(cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184598886859)}
{name auditServiceProviderImpl_1}
{websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditServiceProvider_1173199825608}
{maxFileSize 1}
{websphere_Config_Data_Type AuditServiceProvider}
{fileLocation ${PROFILE_ROOT}/logs/server1}
{className com.ibm.ws.security.audit.BinaryEmitterImpl}
{properties {}}
{eventFormatterClass {}}
{maxLogs 100}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEmitter('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getAuditEmitter(['-emitterRef AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEmitterClass('-interactive')
```

getBinaryFileLocation

The `getBinaryFileLocation` command returns the file location of the binary file audit logs.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a binary file audit service provider implementation. (String, required)

Return value

The command returns the file path of the audit log, as the following example displays:

```
$profile_root/logs/server1
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getBinaryFileLocation('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getBinaryFileLocation(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getBinaryFileLocation('-interactive')
```

getAuditEmitterFilters

The `getAuditEmitterFilters` command returns a list of defined filters for the audit service provider implementation of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies the audit service provider implementation of interest. You can specify a reference to the service provider object. (String, required)

Return value

The command returns a list of defined filters in a shortened format, as the following example output displays:

```
AUTHN:SUCCESS,AUTHN:INFO,AUTHZ:SUCCESS,AUTHZ:INFO
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEmitterFilters('-emitterRef AuditServiceProvider_1173199825608')
```


- Using Jython list:

```
AdminTask.getAuditEmitterFilters(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEmitterFilters('-interactive')
```

getBinaryFileSize

The `getBinaryFileSize` command returns the maximum file size of the binary audit log that is defined for the audit service provider of interest in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a binary file audit service provider implementation. (String, required)

Return value

The command returns the integer value of the maximum file size in megabytes.

Batch mode example usage

- Using Jython string:

```
AdminTask.getBinaryFileSize('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getBinaryFileSize(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getBinaryFileSize('-interactive')
```

getEmitterClass

The `getEmitterClass` command returns the class name of the audit service provider emitter implementation.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a audit service provider implementation. (String, required)

Return value

The command returns the class name of the audit service provider implementation.

Batch mode example usage

- Using Jython string:

```
AdminTask.getEmitterClass('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getEmitterClass(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEmitterClass('-interactive')
```

getEmitterUniqueld

The `getEmitterUniqueld` command returns the unique identifier of the audit service provider implementation.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a service provider implementation. (String, required)

Return value

The command returns the unique ID of the audit service provider of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.getEmitterUniqueId('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getEmitterUniqueId(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEmitterUniqueId('-interactive')
```

getMaxNumBinaryLogs

The `getMaxNumBinaryLogs` command returns the maximum number of binary audit logs that is defined for the audit service provider of interest in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a binary file audit service provider implementation. (String, required)

Return value

The command returns the integer value that represents the maximum number of binary audit logs in the configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.getMaxNumBinaryLogs('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getMaxNumBinaryLogs(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getMaxNumBinaryLogs('-interactive')
```

listAuditEmitters

The listAuditEmitters command returns a list of configured audit service provider implementation objects and the corresponding attributes.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns an array list of audit service provider implementation objects and attributes, as the following example output displays:

```
{auditSpecifications myfilter(cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184598886859)}
{name auditServiceProviderImpl_1}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditServiceProvider_1173199825608}
{maxFileSize 1}
{_Websphere_Config_Data_Type AuditServiceProvider}
{fileLocation ${PROFILE_ROOT}/logs/server1}
{className com.ibm.ws.security.audit.BinaryEmitterImpl}
{properties {}}
{auditSpecRef1 AuditSpecification_1184598886859}
{eventFormatterClass {}}
{maxLogs 100}
{emitterRef AuditServiceProvider_1173199825608}
{{auditSpecifications DefaultAuditSpecification_1(cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1173199825608)}}
{name mythirdpartyemitter}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditServiceProvider_1184686638218}
{maxFileSize 0}
{_Websphere_Config_Data_Type AuditServiceProvider}
{fileLocation {}}
{className com.mycompany.myemitterclass}
{properties {}}
{auditSpecRef1 AuditSpecification_1173199825608}
{eventFormatterClass com.mycompany.myeventformatterclass}
{maxLogs 0}
{emitterRef AuditServiceProvider_1184686638218}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditEmitters()
```

- Using Jython list:

```
AdminTask.listAuditEmitters()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditEmitters('-interactive')
```

modifyAuditEmitter

The `modifyAuditEmitter` command modifies the attributes of an audit service provider implementation object.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a audit service provider implementation. (String, required)

-auditFilters

Specifies a reference or a group of references to predefined audit filters. Use the following format to specify multiple references: `reference,reference,reference` (String, required)

-fileLocation

Specifies the location where the system writes the audit logs. (String, required)

Optional parameters

-className

Specifies the class name to use to identify the implementation. (String, optional)

-eventFormatterClass

Specifies the class that implements how the system formats the audit event for output. If you want to use the default audit service provider, do not specify this parameter. (String, optional)

-customProperties

Specifies a list of custom properties formatted as name and value pairs in the following format: `name=value,name=value`. (String, optional)

-maxFileSize

Specifies the maximum size that each log reaches before the system saves the audit log with a timestamp. Specify the size in megabytes. The default value is 10 MB. (Integer, optional)

-maxLogs

Specifies the maximum number of log files to create before the system rewrites the oldest audit log. The default value is 100 logs. (Integer, optional)

Return value

The command returns a value of `true` if the system successfully modifies the audit service provider of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditEmitter('-emitterRef AuditServiceProvider_1184686638218
-auditFilters AuditSpecification_1173199825608
-fileLocation c:\wasinstall\appserver\profiles\AppSrv01\mylogs -maxFileSize
14 -maxLogs 200')
```

- Using Jython list:

```
AdminTask.modifyAuditEmitter(['-emitterRef', 'AuditServiceProvider_1184686638218',
'-auditFilters', 'AuditSpecification_1173199825608', '-fileLocation',
'c:\wasinstall\appserver\profiles\AppSrv01\mylogs', '-maxFileSize', '14', '-maxLogs',
'200'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditEmitter('-interactive')
```

setAuditEmitterFilters

The setAuditEmitterFilters command sets the filters for an audit service provider implementation.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies a reference to a audit service provider implementation. (String, required)

-filtersRef

Specifies one or more references to defined audit filters. Use the following format to specify more than one filter reference: reference,reference,reference (String, required)

Return value

The command returns a value of true if the system successfully sets the filters for the audit service provider.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAuditEmitterFilters('-emitterRef AuditServiceProvider_1173199825608
-filtersRef AuditSpecification_1184598886859')
```

- Using Jython list:

```
AdminTask.setAuditEmitterFilters(['-emitterRef', 'AuditServiceProvider_1173199825608',
'-filtersRef', 'AuditSpecification_1184598886859'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setAuditEmitterFilters('-interactive')
```

AuditSigningCommands command group for the AdminTask object

You can use the Jython scripting language to configure the signing of audit records with the wsadmin tool. Use the commands and parameters in the AuditSigningCommands group to enable, disable, and configure the security audit system to sign audit records.

Use the following commands to configure audit signing:

- “createAuditSigningConfig” on page 1088
- “deleteAuditSigningConfig” on page 1089
- “disableAuditSigning” on page 1089
- “enableAuditSigning” on page 1090
- “getAuditSigningConfig” on page 1090

- “importEncryptionCertificate” on page 1091
- “isAuditSigningEnabled” on page 1092
- “modifyAuditSigningConfig” on page 1092

createAuditSigningConfig

The createAuditSigningConfig command creates the signing model that the system uses to sign the audit records. Use this command to configure your audit signing configuration for the first time. If you have already configured audit signing, use the enableAuditSigning and disableAuditSigning commands to turn audit signing on and off.

You can import the certificate from an existing key file name containing that certificate, automatically generate the certificate, or use the same certificate as the application server uses to encrypt the audit records. To use an existing certificate in an existing keystore, specify input values for the -enableAuditEncryption, -certAlias, and -signingKeyStoreRef parameters. Also, set the value of the -useEncryptionCert, -autogenCert, and -importCert parameters as false for this scenario.

The user must have the administrator and auditor administrative roles to run this command.

Target object

None.

Required parameters

-enableAuditSigning

Specifies whether to sign audit records. This parameter modifies your audit policy configuration. (Boolean, required)

-certAlias

Specifies the alias name that identifies the generated or imported certificate. (String, required)

-signingKeyStoreRef

Specifies the reference ID of the key store that system imports the certificate to. The signing keystore must already exist in the security.xml file. The system updates this keystore with the certificate that is used to sign the audit records. (String, required)

Optional parameters

-useEncryptionCert

Specifies whether to use the same certificate for encryption and signing. (Boolean, optional)

-autogenCert

Specifies whether to automatically generate the certificate used to sign the audit records. (Boolean, optional)

-importCert

Specifies whether to import an existing certificate to sign the audit records. (Boolean, optional)

-certKeyFileName

Specifies the unique name of the key file for the certificate to import. (String, optional)

-certKeyFilePath

Specifies the key file location for the certificate to import. (String, optional)

-certKeyFileType

Specifies the key file type for the certificate to import. (String, optional)

-certKeyFilePassword

Specifies the key file password for the certificate to import. (String, optional)

-certAliasToImport

Specifies the alias of the certificate to import. (String, optional)

Return value

If successful, returns the shortened form of the keystore where the signing certificate has been added to. Remember, this keystore is in the security.xml file, not the audit.xml file.

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditSigningConfig(['-enableAuditSigning true -certAlias  
auditSigningCert -autogenCert true -signingKeyStoreRef Ref_Id_of_KeyStoreInSecurityXML'])
```

- Using Jython list:

```
AdminTask.createAuditSigningConfig(['-enableAuditSigning', 'true', '-certAlias',  
'auditSigningCert', '-autogenCert', 'true -signingKeyStoreRef',  
'Ref_Id_of_KeyStoreInSecurityXML'])
```

Interactive mode example usage

- Using Jython :

```
AdminTask.createAuditSigningConfig('-interactive')
```

deleteAuditSigningConfig

The deleteAuditSigningConfig command deletes the signing model that the system uses to sign the audit records. When the system deletes the audit signing configuration, it does not delete the key store file in the security.xml or the signer certificate for the keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of true if the system successfully removes the audit signing configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditSigningConfig()
```

- Using Jython list:

```
AdminTask.deleteAuditSigningConfig()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditSigningConfig('-interactive')
```

disableAuditSigning

The disableAuditSigning command disables audit record signing for the security auditing system.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully disables audit signing.

Batch mode example usage

- Using Jython string:

```
AdminTask.disableAuditSigning()
```

- Using Jython list:

```
AdminTask.disableAuditSigning()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.disableAuditSigning('-interactive')
```

enableAuditSigning

The `enableAuditSigning` command enables audit record signing in the security auditing system.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully enables audit signing in the security auditing system.

Batch mode example usage

- Using Jython string:

```
AdminTask.enableAuditSigning()
```

- Using Jython list:

```
AdminTask.enableAuditSigning()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.enableAuditSigning()
```

getAuditSigningConfig

The `getAuditSigningConfig` command retrieves the signing model that the system uses to sign the audit records.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes that are associated with the signing model, as the following sample output displays:

```
{{securityXmlSignerScopeName (cell):Node04Cell:(node):Node04}
{securityXmlSignerCertAlias mysigningcert}
{securityXmlSignerKeyStoreName NodeDefaultRootStore}
{signerKeyStoreRef KeyStore_Node04_4}
{enabled true}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditSigningConfig()
```

- Using Jython list:

```
AdminTask.getAuditSigningConfig()
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.getAuditSigningConfig('-interactive')
```

importEncryptionCertificate

The `importEncryptionCertificate` command imports the self-signed certificate used for encrypting audit data from the encryption keystore into another keystore. Use this command internally to automatically generate a certificate for either encryption or signing. You can also use this command to import the certificate into the keystore by specifying the `keyStoreName` and `keyStoreScope` parameters.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name to identify the keystore. (String, required)

-keyFilePath

Specifies the keystore path name that contains the certificate to import. (String, required)

-keyFilePassword

Specifies the password of the keystore that contains the certificate to import. (String, required)

-keyFileType

Specifies the type of the keystore. (String, required)

-certificateAliasFromKeyFile

Specifies the alias of the certificate to import from the keystore file. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-certificateAlias

Specifies a unique name to identify the imported certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully imports the encryption certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.importEncryptionCertificate('{-keyStoreName AuditDefaultKeyStore -keyStoreScope
(cell):Node04Cell -keyFilePath c:/install_root/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12
-keyFilePassword WebAS -keyFileType PKCS12 -certificateAliasFromKeyFile root -certificateAlias myimportcert'})
```

- Using Jython list:

```
AdminTask.importEncryptionCertificate(['-keyStoreName', 'AuditDefaultKeyStore', '-keyStoreScope',  
'(cell):Node04Cell', '-keyFilePath', 'c:/install_root/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12',  
'-keyFilePassword', 'WebAS', '-keyFileType', 'PKCS12', '-certificateAliasFromKeyFile',  
'root', '-certificateAlias', 'myimportcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.importEncryptionCertificate('-interactive')
```

isAuditSigningEnabled

The `isAuditSigningEnabled` command indicates whether audit signing is enabled or disabled in the security audit system.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if signing is configured in the security auditing system.

Batch mode example usage

- Using Jython string:

```
AdminTask.isAuditSigningEnabled()
```

- Using Jython list:

```
AdminTask.isAuditSigningEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isAuditSigningEnabled('-interactive')
```

modifyAuditSigningConfig

The `modifyAuditSigningConfig` command modifies the signing model that the system uses to sign the audit records.

The certificate may either be imported from an existing key file name containing that certificate, automatically generated, or be the same certificate used to encrypt the audit records. To use an existing certificate in an existing keystore, specify input values for the `-enableAuditEncryption`, `-certAlias`, and `-signingKeyStoreRef` parameters. Also, set the value the `-useEncryptionCert`, `-autogenCert`, and `-importCert` parameters as `false` for this scenario.

The user must have the administrator and auditor administrative roles to run this command.

Target object

None.

Required parameters

-enableAuditSigning

Specifies whether to sign audit records. This parameter modifies your audit policy configuration. (Boolean, required)

-certAlias

Specifies the alias name that identifies the generated or imported certificate. (String, required)

-signingKeyStoreRef

Specifies the reference ID of the key store that system imports the certificate to. The signing keystore must already exist in the security.xml file. The system updates this keystore with the certificate that is used to sign the audit records. (String, required)

Optional parameters

-useEncryptionCert

Specifies whether to use the same certificate for encryption and signing. (Boolean, optional)

-autogenCert

Specifies whether to automatically generate the certificate used to sign the audit records. (Boolean, optional)

-importCert

Specifies whether to import an existing certificate to sign the audit records. (Boolean, optional)

-certKeyFileName

Specifies the unique name of the key file for the certificate to import. (String, optional)

-certKeyFilePath

Specifies the key file location for the certificate to import. (String, optional)

-certKeyFileType

Specifies the key file type for the certificate to import. (String, optional)

-certKeyFilePassword

Specifies the key file password for the certificate to import. (String, optional)

-certAliasToImport

Specifies the alias of the certificate to import. (String, optional)

Return value

The command returns a value of `true` if the system successfully modifies the security auditing system configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditSigningConfig('-enableAuditSigning true -certAlias auditSigningCert
  -autogenCert true -signingKeyStoreRef Ref_Id_of_KeyStoreInSecurityXML')
```

- Using Jython list:

```
AdminTask.modifyAuditSigningConfig(['-enableAuditSigning', 'true', '-certAlias',
  'auditSigningCert', '-autogenCert', 'true', '-signingKeyStoreRef', 'Ref_Id_of_KeyStoreInSecurityXML'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditSigningConfig('-interactive')
```

AuditEncryptionCommands command group for the AdminTask object

You can use the Jython scripting language to configure the security auditing system with the `wsadmin` tool. Use the commands and parameters in the `AuditEncryptionCommands` group to configure the security audit system to encrypt audit records.

Use the following commands to enable, disable, and configure audit record encryption:

- “`createAuditEncryptionConfig`” on page 1094

- “createAuditSelfSignedCertificate” on page 1095
- “deleteAuditCertificate” on page 1096
- “deleteAuditEncryptionConfig” on page 1097
- “disableAuditEncryption” on page 1098
- “enableAuditEncryption” on page 1098
- “exportAuditCertificate” on page 1098
- “exportAuditCertToManagedKS” on page 1099
- “getAuditCertificate” on page 1100
- “getAuditEncryptionConfig” on page 1101
- “getEncryptionKeyStore” on page 1102
- “importAuditCertFromManagedKS” on page 1102
- “importAuditCertificate” on page 1103
- “importEncryptionCertificate” on page 1104
- “isAuditEncryptionEnabled” on page 1105
- “listAuditEncryptionKeyStores” on page 1106
- “listCertAliases” on page 1106
- “modifyAuditEncryptionConfig” on page 1107
- “renewAuditCertificate” on page 1108

createAuditEncryptionConfig

The createAuditEncryptionConfig command creates the encryption model used to encrypt the audit records.

You can import the certificate from an existing key file name containing that certificate or automatically generate a certificate.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-enableAuditEncryption

Specifies whether to encrypt audit records. This parameter modifies your audit policy configuration. (Boolean, required)

-certAlias

Specifies the alias name that identifies the generated or imported certificate. (String, required)

-encryptionKeyStoreRef

Specifies the reference ID of the keystore to import the certificate to. (String, required)

Optional parameters

-autogenCert

Specifies whether to automatically generate the certificate used to encrypt the audit records. You must specify either this parameter or the -importCert parameter, but you cannot specify both. (Boolean, optional)

-importCert

Specifies whether to import an existing certificate to encrypt the audit records. You must specify either this parameter or the `-autogenCert` parameter, but you cannot specify both. (Boolean, optional)

-certKeyFileName

Specifies the unique name of the key file for the certificate to import. (String, optional)

-certKeyFilePath

Specifies the key file location for the certificate to import. (String, optional)

-certKeyFileType

Specifies the key file type for the certificate to import. (String, optional)

-certKeyFilePassword

Specifies the key file password for the certificate to import. (String, optional)

-certAliasToImport

Specifies the alias of the certificate to import. (String, optional)

Return value

The command returns the shortened form of the reference ID of the created encryption keystore if the system successfully creates the audit encryption configuration, as the following example output displays:

```
KeyStore_1173199825578
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditEncryptionConfig('-enableAuditEncryption true -certAlias  
auditCertificate -autogenCert true -encryptionKeyStoreRef auditKeyStore')
```

- Using Jython list:

```
AdminTask.createAuditEncryptionConfig(['-enableAuditEncryption', 'true', '-certAlias',  
'auditCertificate', '-autogenCert', 'true', '-encryptionKeyStoreRef', 'auditKeyStore'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.createAuditEncryptionConfig('-interactive')
```

createAuditSelfSignedCertificate

The `createAuditSelfSignedCertificate` command creates a self-signed certificate. Use this command internally to automatically generate a certificate for encryption and signing or to import that certificate into the keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore where the system imports the self-signed certificate to. (String, optional)

-certificateAlias

Specifies a unique alias name for the certificate. (String, required)

-certificateSize

Specifies the size that the private key uses for the personal certificate. The default value is 1024. (Integer, required)

-certificateCommonName

Specifies the common name portion of the distinguished name. (String, required)

Optional parameters

-certificateOrganization

Specifies the organizational part of the distinguished name. (String, optional)

-keyStoreScope

Specifies the scope of the keystore that the system imports the self-signed certificate to. (String, optional)

-certificateVersion

Specifies the version of the personal certificate. (String, optional)

-certificateOrganizationalUnit

Specifies the organization unit part of the distinguished name. (String, optional)

-certificateLocality

Specifies the locality portion of the distinguished name. (String, optional)

-certificateState

Specifies the state portion of the distinguished name. (String, optional)

-certificateZip

Specifies the zip code portion of the distinguished name. (String, optional)

-certificateCountry

Specifies the country portion of the distinguished name. The default value is US. (String, optional)

-certificateValidDays

Specifies the length of time, in days, which the certificate is valid. The default value is 365 days. (Integer, optional)

Return value

The command returns a value of `true` if the system successfully creates the self-signed certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditSelfSignedCertificate('-keyStoreName AuditDefaultKeyStore -keyStoreScope  
(cell):Node04Cell -certificateAlias myNew -certificateCommonName cn=oet -certificateOrganization mycompany')
```

- Using Jython list:

```
AdminTask.createAuditSelfSignedCertificate(['-keyStoreName', 'AuditDefaultKeyStore', '-keyStoreScope',  
'(cell):Node04Cell', '-certificateAlias', 'myNew', '-certificateCommonName', 'cn=oet',  
'-certificateOrganization', 'mycompany'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditSelfSignedCertificate('-interactive')
```

deleteAuditCertificate

The `deleteAuditCertificate` command deletes a self-signed certificate from an audit keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore from which the system deletes the self-signed certificate. (String, required)

-certificateAlias

Specifies a unique alias name for the certificate to delete. (String, required)

Optional parameters

-keyStoreScope

Specifies a unique alias name for the certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully deletes the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditCertificate('-keyStoreName myKeystore -certificateAlias oldCertificate')
```

- Using Jython list:

```
AdminTask.deleteAuditCertificate(['-keyStoreName', 'myKeystore', '-certificateAlias', 'oldCertificate'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditCertificate('-interactive')
```

deleteAuditEncryptionConfig

The `deleteAuditEncryptionConfig` command deletes the encryption model used to encrypt the audit records. The command does not remove keystore files or the certificates.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully deletes the audit encryption configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditEncryptionConfig()
```

- Using Jython list:

```
AdminTask.deleteAuditEncryptionConfig()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditEncryptionConfig('-interactive')
```

disableAuditEncryption

The `disableAuditEncryption` command disables the encryption of audit records.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully disables audit record encryption.

Batch mode example usage

- Using Jython string:

```
AdminTask.disableAuditEncryption()
```

- Using Jython list:

```
AdminTask.disableAuditEncryption()
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.disableAuditEncryption('-interactive')
```

enableAuditEncryption

The `enableAuditEncryption` command enables the encryption of audit records.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully enables audit record encryption.

Batch mode example usage

- Using Jython string:

```
AdminTask.enableAuditEncryption()
```

- Using Jython list:

```
AdminTask.enableAuditEncryption()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.enableAuditEncryption()
```

exportAuditCertificate

The `exportAuditCertificate` command exports a self-signed certificate from a keystore. To use this command, you must adhere to the following user role and privilege guidelines:

- You must have audit privileges to export the certificate from an audit keystore.
- You must have the auditor and administrator roles to export the certificate to a security keystore.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

-keyStorePassword

Specifies the password that the system uses to access the keystore specified with the `-keyStoreName` parameter. (String, required)

-keyFilePath

Specifies the key store path name that contains the certificate to export. (String, required)

-keyFilePassword

Specifies the password of the keystore that contains the certificate to export. (String, required)

-keyFileType

Specifies the type of the keystore. (String, required)

-certificateAlias

Specifies the alias of the certificate to export from the keystore. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-aliasInKeyStore

Specifies a new unique name to identify the exported certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully exports the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.exportAuditCertificate('-keyStoreName AuditDefaultKeyStore -keyStoreScope  
(cell):Node04Cell -keyFilePath c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12  
-keyFilePassword myPwd -keyFileType PKCS12 -certificateAlias root')
```

- Using Jython list:

```
AdminTask.exportAuditCertificate(['-keyStoreName', 'AuditDefaultKeyStore', '-keyStoreScope',  
'(cell):Node04Cell', '-keyFilePath', 'c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12',  
-keyFilePassword', 'myPwd', '-keyFileType', 'PKCS12', '-certificateAlias', 'root'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.exportAuditCertificate('-interactive')
```

exportAuditCertToManagedKS

The `exportAuditCertToManagedKS` command exports a self-signed certificate from an audit keystore to a managed audit keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the managed keystore. (String, required)

-keyStorePassword

Specifies the password of the managed keystore that contains the certificate to export. (String, required)

-toKeyStoreName

Specifies the unique name of the managed keystore that contains the certificate to export. (String, required)

-certificateAlias

Specifies a unique name to identify the exported certificate. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-toKeyStoreScope

Specifies the scope of the managed keystore that contains the certificate to export. (String, optional)

-aliasInKeyStore

Specifies the new unique name to identify the exported certificate. If you do not specify a value for this parameter, the system sets the unique name to the value specified for the `-certificateAlias` parameter. (String, optional)

Return value

The command returns a value of `true` if the system successfully exports the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.exportAuditCertToManagedKS('-keyStoreName auditEncryptionKeyStore -keyStorePassword myPwd  
-toKeyStoreName AuditTrustStore -toKeyStoreScope (cell):my03Cell -certificateAlias newauditcert  
-aliasInKeyStore newauditcert1')
```

- Using Jython list:

```
AdminTask.exportAuditCertToManagedKS(['-keyStoreName', 'auditEncryptionKeyStore', '-keyStorePassword', 'myPwd',  
'-toKeyStoreName', 'AuditTrustStore', '-toKeyStoreScope', '(cell):my03Cell', '-certificateAlias', 'newauditcert',  
'-aliasInKeyStore', 'newauditcert1'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.exportAuditCertToManagedKS('-interactive')
```

getAuditCertificate

The `getAuditCertificate` command retrieves the attributes for an audit self-signed certificate in an audit keystore.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the managed keystore of interest. (String, required)

-certificateAlias

Specifies a unique name to identify the exported certificate of interest. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore of interest. (String, optional)

Return value

The command returns a list of attributes associated with the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditCertificate('-keyStoreName auditEncryptionKeyStore -certificateAlias
newauditcert')
```

- Using Jython list:

```
AdminTask.getAuditCertificate(['-keyStoreName', 'auditEncryptionKeyStore', '-certificateAlias',
'newauditcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditCertificate('-interactive')
```

getAuditEncryptionConfig

The `getAuditEncryptionConfig` command retrieves the encryption model that the system uses to encrypt the audit records.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes associated with the encryption model, as the following example output displays:

```
{{certRef Certificate_1184698729015}
{keystoreRef KeyStore_1173199825578}
{keyStore AuditDefaultKeyStore(cells/CHEYENNENode04Cell|audit.xml#KeyStore_1173199825578)}
{enabled true}
{alias mycertalias}
{ _Websphere_Config_Data_Version {} }
{ _Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#Certificate_1184698729015}
{ _Websphere_Config_Data_Type Certificate}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEncryptionConfig()
```

- Using Jython list:

```
AdminTask.getAuditEncryptionConfig()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEncryptionConfig('-interactive')
```

getEncryptionKeyStore

The `getEncryptionKeyStore` command retrieves the attributes for the keystore that contains the certificate that the system uses to encrypt the audit records.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes for the keystore of interest, as the following example displays:

```
{location ${CONFIG_ROOT}/audittrust.p12}
{password *****}
{ Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#KeyStore_1173199825578}
{ Websphere_Config_Data_Version {} }
{useForAcceleration false}
{slot 0}
{type PKCS12}
{additionalKeyStoreAttrs {}}
{fileBased true}
{ Websphere_Config_Data_Type KeyStore}
{customProviderClass {}}
{hostList {}}
{keystoreRef KeyStore_1173199825578}
{createStashFileForCMS false}
{description {keyStore description}}
{managementScope (cells/CHEYENNENode04Cell|audit.xml#ManagementScope_1173199825608)}
{readOnly false}
{initializeAtStartup true}
{usage {}}
{provider IBMJCE}
{name AuditDefaultKeyStore}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getEncryptionKeyStore()
```

- Using Jython list:

```
AdminTask.getEncryptionKeyStore()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEncryptionKeyStore('-interactive')
```

importAuditCertFromManagedKS

The `importAuditCertFromManagedKS` command imports a self-signed certificate into a keystore from a managed audit keystore. Use this command internally to automatically generate a certificate for encryption or signing and to import a certificate into the keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the managed keystore. (String, required)

-fromKeyStoreName

Specifies the unique name of the managed keystore that contains the certificate to import. (String, required)

-fromKeyStorePassword

Specifies the password of the managed keystore that contains the certificate to import. (String, required)

-certificateAliasFromKeyFile

Specifies the alias of the certificate to import from the managed keystore file. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-fromKeyStoreScope

Specifies the scope of the managed keystore that contains the certificate to import. (String, optional)

-certificateAlias

Specifies a unique name to identify the imported certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully imports the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.importAuditCertFromManagedKS('-keyStoreName AuditDefaultKeyStore -keyStoreScope
(cell):myNode03Cell -fromKeyStoreName AuditSecondDefaultKeyStore -fromKeyStoreScope
(cell):myNode03Cell -fromKeyStorePassword myPwd
-certificateAliasFromKeyFile root -certificateAlias myimportcert')
```

- Using Jython list:

```
AdminTask.importAuditCertFromManagedKS(['-keyStoreName', 'AuditDefaultKeyStore', '-keyStoreScope',
'(cell):Node04Cell', '-fromKeyStoreName', 'AuditSecondDefaultKeyStore', '-fromKeyStoreScope',
'(cell):myNode03Cell', '-fromKeyStorePassword', 'myPwd', '-certificateAliasFromKeyFile',
'root', '-certificateAlias', 'myimportcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.importAuditCertFromManagedKS('-interactive')
```

importAuditCertificate

The `importAuditCertificate` command imports a self-signed certificate into a keystore. Use this command internally to automatically generate a certificate for encryption or signing and to import a certificate into the keystore. To use this command, you must adhere to the following user role and privilege guidelines:

- You must have audit privileges to import the certificate to an audit keystore.
- You must have the auditor and administrator roles to import the certificate to a security keystore.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

-keyFilePath

Specifies the key store path name that contains the certificate to import. (String, required)

-keyFilePassword

Specifies the password of the keystore that contains the certificate to import. (String, required)

-keyFileType

Specifies the type of the keystore. (String, required)

-certificateAliasFromKeyFile

Specifies the alias of the certificate to import from the keystore file. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-certificateAlias

Specifies a unique name to identify the imported certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully imports the audit certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.importAuditCertificate('-keyStoreName AuditDefaultKeyStore -keyStoreScope  
(cell):Node04Cell -keyFilePath c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12  
-keyFilePassword myPwd -keyFileType PKCS12 -certificateAliasFromKeyFile root -certificateAlias myimportcert')
```

- Using Jython list:

```
AdminTask.importAuditCertificate(['-keyStoreName', 'AuditDefaultKeyStore', '-keyStoreScope', '(cell):Node04Cell',  
'-keyFilePath', 'c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12',  
'-keyFilePassword', 'myPwd', '-keyFileType', 'PKCS12', '-certificateAliasFromKeyFile', 'root',  
'-certificateAlias', 'myimportcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.importAuditCertificate('-interactive')
```

importEncryptionCertificate

The `importEncryptionCertificate` command imports the self-signed certificate that the system uses to encrypt audit data from the encryption keystore into a managed keystore in `security.xml`.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

-keyFilePath

Specifies the key store path name that contains the certificate to import. (String, required)

-keyFilePassword

Specifies the password of the keystore that contains the certificate to import. (String, required)

-keyFileType

Specifies the type of the keystore. (String, required)

-certificateAliasFromKeyFile

Specifies the alias of the certificate to import from the keystore file. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore. (String, optional)

-certificateAlias

Specifies a unique name to identify the imported certificate. (String, optional)

Return value

The command returns a value of `true` if the system successfully imports the encryption certificate.

Batch mode example usage

- Using Jython string:

```
AdminTask.importEncryptionCertificate('-keyStoreName DefaultKeyStore -keyStoreScope (cell):Node04Cell  
-keyFilePath c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12  
-keyFilePassword myPwd -keyFileType PKCS12 -certificateAliasFromKeyFile root -certificateAlias myimportcert')
```

- Using Jython list:

```
AdminTask.importEncryptionCertificate(['-keyStoreName', 'DefaultKeyStore', '-keyStoreScope', '(cell):Node04Cell',  
'-keyFilePath', 'c:/wasinstall/appserver/profiles/AppSrv01/config/cells/Node04Cell/nodes/Node04/trust.p12',  
'-keyFilePassword', 'myPwd', '-keyFileType', 'PKCS12', '-certificateAliasFromKeyFile', 'root',  
'-certificateAlias', 'myimportcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.importEncryptionCertificate('-interactive')
```

isAuditEncryptionEnabled

The `isAuditEncryptionEnabled` command determines if audit record encryption is enabled.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if audit record encryption is enabled.

Batch mode example usage

- Using Jython string:

```
AdminTask.isAuditEncryptionEnabled()
```

- Using Jython list:

```
AdminTask.isAuditEncryptionEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isAuditEncryptionEnabled('-interactive')
```

listAuditEncryptionKeyStores

The `listAuditEncryptionKeyStores` command retrieves the attributes for each configured encryption keystore from the `audit.xml` file. The command returns attributes for active and inactive keystores.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes for each configured keystore, as the following example output displays:

```
{{location ${CONFIG_ROOT}/audittrust.p12}
{password *****}
{websphere_config_data_id cells/CHEYENNENode04Cell|audit.xml#KeyStore_1173199825578}
{useForAcceleration false}
{slot 0}
{type PKCS12}
{additionalKeyStoreAttrs {}}
{fileBased true}
{websphere_config_data_type KeyStore}
{customProviderClass {}}
{hostList {}}
{keystoreRef KeyStore_1173199825578}
{createStashFileForCMS false}
{description {keyStore description}}
{readOnly false}
{initializeAtStartup true}
{managementScope (cells/CHEYENNENode04Cell|audit.xml#ManagementScope_1173199825608)}
{usage {}}
{provider IBMJCE}
{name AuditDefaultKeyStore}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditEncryptionKeyStores()
```

- Using Jython list:

```
AdminTask.listAuditEncryptionKeyStores()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditEncryptionKeyStores('-interactive')
```

listCertAliases

The `listCertAliases` command retrieves a list of the personal certificates in the keystore, as specified by the keystore name and scope of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-keyStoreName

Specifies the unique name of the keystore. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope of the keystore. The default value is the cell scope. (String, optional)

Return value

The command returns a list of certificate aliases for the personal certificates that are configured for the keystore, as the following sample output displays:

```
mycertalias
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listCertAliases('-keyStoreName AuditDefaultKeyStore -keyStoreScope (cell):Node04Cell')
```

- Using Jython list:

```
AdminTask.listCertAliases(['-keyStoreName AuditDefaultKeyStore -keyStoreScope (cell):Node04Cell'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listCertAliases('-interactive')
```

modifyAuditEncryptionConfig

The `modifyAuditEncryptionConfig` command modifies the encryption model that the system uses to encrypt the audit records. Specify values for the `-enableAuditEncryption`, `-certAlias`, and `encryptionKeyStoreRef` parameters to use an existing keystore. Do not specify the `-importCert` or `-autogenCert` parameters if you use an existing keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

None

Optional parameters

-enableAuditEncryption

Specifies whether to encrypt audit records. This parameter modifies your audit policy configuration. (Boolean, optional)

-autogenCert

Specifies whether to automatically generate the certificate used to encrypt the audit records. You must specify either this parameter or the `-importCert` parameter, but you cannot specify both. (Boolean, optional)

-importCert

Specifies whether to import an existing certificate to encrypt the audit records. You must specify either this parameter or the `-autogenCert` parameter, but you cannot specify both. (Boolean, optional)

-certKeyFileName

Specifies the unique name of the key file for the certificate to import. (String, optional)

-certKeyFilePath

Specifies the key file location for the certificate to import. (String, optional)

-certKeyFileType

Specifies the key file type for the certificate to import. (String, optional)

-certKeyFilePassword

Specifies the key file password for the certificate to import. (String, optional)

-certAliasToImport

Specifies the alias of the certificate to import. (String, optional)

-certAlias

Specifies the alias name that identifies the generated or imported certificate. (String, optional)

-encryptionKeyStoreRef

Specifies the reference ID of the keystore to import the certificate to. (String, optional)

Return value

The command returns a value of `true` if the system successfully updates the configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditEncryptionConfig('-enableAuditEncryption true -certAlias mycertalias  
-encryptionKeyStoreRef KeyStore_1173199825578')
```

- Using Jython list:

```
AdminTask.modifyAuditEncryptionConfig(['-enableAuditEncryption', 'true', '-certAlias', 'mycertalias',  
'-encryptionKeyStoreRef', 'KeyStore_1173199825578'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditEncryptionConfig('-interactive')
```

renewAuditCertificate

The `renewAuditCertificate` command renews a self signed certificate in an audit keystore.

The user must have the auditor administrative role to run this command.

Target object

None.

-keyStoreName

Specifies the unique name of the managed keystore of interest. (String, required)

-certificateAlias

Specifies a unique name to identify the exported certificate to renew. (String, required)

Optional parameters

-keyStoreScope

Specifies the scope name of the keystore of interest. (String, optional)

Return value

The command returns a value of `true` if the system successfully updates the configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.renewAuditCertificate('-keyStoreName auditEncryptionKeyStore  
-certificateAlias newauditcert')
```

- Using Jython list:

```
AdminTask.renewAuditCertificate(['-keyStoreName', 'auditEncryptionKeyStore',
'-certificateAlias', 'newauditcert'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.renewAuditCertificate('-interactive')
```

AuditEventFactoryCommands for the AdminTask object

You can use the Jython scripting language to configure the security auditing system with the wsadmin tool. Use the commands and parameters in the AuditEventFactoryCommands group to configure the default or a third-party audit event factory.

Use the following commands to configure the default audit event factory or a third-party audit event factory:

- “createAuditEventFactory”
- “deleteAuditEventFactoryByName” on page 1110
- “deleteAuditEventFactoryByRef” on page 1111
- “getAuditEventFactory” on page 1111
- “getAuditEventFactoryClass” on page 1112
- “getAuditEventFactoryFilters” on page 1112
- “getAuditEventFactoryName” on page 1113
- “getAuditEventFactoryProvider” on page 1114
- “listAuditEventFactories” on page 1114
- “modifyAuditEventFactory” on page 1115
- “setAuditEventFactoryFilters” on page 1116

createAuditEventFactory

The createAuditEventFactory command creates an audit event factory in your security auditing system configuration. You can use the default implementation of the audit event factory or use a third-party implementation. To configure a third-party implementation, use the optional -customProperties parameter to specify any properties necessary to configure the audit event factory implementation.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies a unique name that identifies the audit event factory. (String, required)

-className

Specifies the class implementation of the audit event factory interface. (String, required)

-provider

Specifies a reference to a predefined audit service provider implementation. (String, required)

-auditFilters

Specifies a reference or a group of references to predefined audit filters, using the following format: reference, reference, reference (String, required)

Optional parameters

-customProperties

Specifies any custom properties necessary to configure a third-party implementation. (String, optional)

Return value

The command returns the shortened reference ID for the newly created audit event factory.

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditEventFactory('-uniqueName myeventfactory -className  
com.mycompany.myeventfactoryclass -provider AuditServiceProvider_1173199825608  
-customProperties a=b -auditFilters AuditSpecification_1184598886859')
```

- Using Jython list:

```
AdminTask.createAuditEventFactory(['-uniqueName', 'myeventfactory', '-className',  
'com.mycompany.myeventfactoryclass', '-provider', 'AuditServiceProvider_1173199825608',  
'-customProperties', 'a=b', '-auditFilters', 'AuditSpecification_1184598886859'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditEventFactory()
```

deleteAuditEventFactoryByName

The `deleteAuditEventFactoryByName` command deletes the audit event factory implementation in the `audit.xml` file that matches a specific unique name identifier.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-uniqueName

Specifies the unique name of the audit event factory implementation. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit event factory.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditEventFactoryByName('-uniqueName  
myeventfactory')
```

- Using Jython list:

```
AdminTask.deleteAuditEventFactoryByName(['-uniqueName', 'myeventfactory'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditEventFactoryByName('-interactive')
```


deleteAuditEventFactoryByRef

The `deleteAuditEventFactoryByRef` command deletes the audit event factory implementation that matches the reference ID of interest.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit event factory.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditEventFactoryByRef('-eventFactoryRef  
AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.deleteAuditEventFactoryByRef(['-eventFactoryRef', 'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditEventFactoryByRef('-interactive')
```

getAuditEventFactory

The `getAuditEventFactory` command retrieves the list of attributes for the audit event factory implementation in the `audit.xml` file for a specific reference id.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns an attribute list for the audit event factory implementation of interest, as the following example output displays:

```
{{name myeventfactory}  
{properties {{{validationExpression {}}}  
{name a}  
{description {}}  
{value b}}
```

```
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#Property_1184688593531}
{_Websphere_Config_Data_Type Property}
{required false}}}}
{className com.mycompany.myeventfactoryclass}
{auditServiceProvider auditServiceProviderImpl_1(cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608)}
{auditSpecifications DefaultAuditSpecification_1(cells/Node04Cell|audit.xml#AuditSpecification_1173199825608)}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditEventFactory_1184688293515}
{_Websphere_Config_Data_Type AuditEventFactory}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEventFactory('-eventFactoryRef AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.getAuditEventFactory(['-eventFactoryRef', 'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEventFactory('-interactive')
```

getAuditEventFactoryClass

The `getAuditEventFactoryClass` command retrieves the class name of the audit event factory implementation that matches a specific reference ID in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns the class name of the audit event factory of interest, as the following sample output displays:

```
com.mycompany.myeventfactoryclass
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEventFactoryClass('-eventFactoryRef
AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.getAuditEventFactoryClass(['-eventFactoryRef',
'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEventFactoryClass('-interactive')
```

getAuditEventFactoryFilters

The `getAuditEventFactoryFilters` command retrieves a list of defined filters for the passed-in event factory.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns a list of the defined filters for the event factory reference of interest in a shortened format, as the following sample output displays:

```
AUTHN:SUCCESS,AUTHN:INFO,AUTHZ:SUCCESS,AUTHZ:INFO
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEventFactoryFilters('-eventFactoryRef  
AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.getAuditEventFactoryFilters(['-eventFactoryRef',  
'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEventFactoryFilters('-interactive')
```

getAuditEventFactoryName

The `getAuditEventFactoryName` command retrieves the unique name of the audit event factory implementation that matches a specific reference ID in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns the name of the audit event factory, as the following sample output displays:

```
myeventfactory
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEventFactoryName('-eventFactoryRef  
AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.getAuditEventFactoryName(['-eventFactoryRef',  
'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEventFactoryName('-interactive')
```

getAuditEventFactoryProvider

The `getAuditEventFactoryProvider` command retrieves the object name of the audit service provider that a specific audit event factory implementation uses in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Return value

The command returns the object name of the audit service provider for the audit event factory of interest, as the following sample output displays:

```
auditServiceProviderImp_1(cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608)
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditEventFactoryProvider('-eventFactoryRef  
AuditEventFactory_1184688293515')
```

- Using Jython list:

```
AdminTask.getAuditEventFactoryProvider(['-eventFactoryRef',  
'AuditEventFactory_1184688293515'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditEventFactoryProvider('-interactive')
```

listAuditEventFactories

The `listAuditEventFactories` command retrieves a list of audit event factory objects and their attributes that are defined in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns an array list of audit event factories and attributes, as the following example output displays:

```
{{auditSpecifications DefaultAuditSpecification_1(cells/Node04Cell|audit.xml#AuditSpecification_1173199825608)  
DefaultAuditSpecification_2(cells/Node04Cell|audit.xml#AuditSpecification_1173199825609)  
DefaultAuditSpecification_3(cells/Node04Cell|audit.xml#AuditSpecification_1173199825610)
```

```

DefaultAuditSpecification_4(cells/Node04Cell|audit.xml#AuditSpecification_1173199825611)}
{name auditEventFactoryImpl_1}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditEventFactory_1173199825608}
{_Websphere_Config_Data_Type AuditEventFactory}
{auditSpecRef4 AuditSpecification_1173199825611}
{properties {}}
{auditSpecRef3 AuditSpecification_1173199825610}
{className com.ibm.ws.security.audit.AuditEventFactoryImpl}
{auditServiceProvider auditServiceProviderImpl_1(cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608)}
{auditSpecRef2 AuditSpecification_1173199825609}
{auditSpecRef1 AuditSpecification_1173199825608}
{auditEventFactoryRef AuditEventFactory_1173199825608}
{emitterRef AuditServiceProvider_1173199825608)}
{{auditSpecifications myfilter(cells/Node04Cell|audit.xml#AuditSpecification_1184598886859)}}
{name myeventfactory}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditEventFactory_1184688293515}
{_Websphere_Config_Data_Type AuditEventFactory}
{className com.mycompany.myeventfactoryclass}
{auditServiceProvider auditServiceProviderImpl_1(cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608)}
{properties {{{validationExpression {}}
{name a}
{description {}}
{value b}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#Property_1184688293546}
{_Websphere_Config_Data_Type Property}
{required false}}}}}
{auditSpecRef1 AuditSpecification_1184598886859}
{auditEventFactoryRef AuditEventFactory_1184688293515}
{emitterRef AuditServiceProvider_1173199825608)}

```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditEventFactories()
```

- Using Jython list:

```
AdminTask.listAuditEventFactories()
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.listAuditEventFactories('-interactive')
```

modifyAuditEventFactory

The modifyAuditEventFactory command modifies the attributes of the audit event factory implementation that the command references with the reference id.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

Optional parameters

-provider

Specifies the reference ID of an audit service provider implementation. (String, optional)

-className

Specifies the name of the class that implements the audit event factory interface. (String, optional)

-customProperties

Specifies one or more custom properties to associate with the audit event factory of interest. Use the following format: name=value, name=value (String, optional)

-auditFilters

Specifies a list of references to audit filters that exist in your configuration. You can separate each item in the list with a comma (,), a semicolon (;), or a space. (String, optional)

Return value

The command returns a value of `true` if the system successfully updates the security auditing system configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditEventFactory('-eventFactoryRef AuditEventFactory_1184688293515 -provider AuditServiceProvider_1173199825608 -customProperties b=c')
```

- Using Jython list:

```
AdminTask.modifyAuditEventFactory(['-eventFactoryRef', 'AuditEventFactory_1184688293515', '-provider', 'AuditServiceProvider_1173199825608', '-customProperties', 'b=c'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditEventFactory('-interactive')
```

setAuditEventFactoryFilters

The `setAuditEventFactoryFilters` command sets the filters for an audit event factory implementation.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-eventFactoryRef

Specifies an audit event factory implementation. This parameter can be a reference to the event factory object. (String, required)

-filtersRef

Specifies a list of references to defined audit filters. (String, required)

Return value

The command returns a value of `true` if the system successfully sets the filters for the audit event factory.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAuditEventFactoryFilters('-eventFactoryRef AuditEventFactory_1184688293515 -filtersRef AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.setAuditEventFactoryFilters(['-eventFactoryRef', 'AuditEventFactory_1184688293515', '-filtersRef', 'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.setAuditEventFactoryFilters('-interactive')
```

AuditFilterCommands command group for the AdminTask object

You can use the Jython scripting language to configure the security auditing system with the wsadmin tool. Use the commands and parameters in the AuditFilterCommands group to configure and manage auditable events.

Use the following commands to configure filters for auditable events in your security auditing configuration:

- “convertFilterRefToString”
- “convertFilterStringToRef” on page 1118
- “createAuditFilter” on page 1118
- “deleteAuditFilter” on page 1120
- “deleteAuditFilterByRef” on page 1121
- “disableAuditFilter” on page 1121
- “enableAuditFilter” on page 1122
- “getAuditFilter” on page 1122
- “getAuditOutcomes” on page 1123
- “getSupportedAuditEvents” on page 1124
- “getSupportedAuditOutcomes” on page 1125
- “isAuditFilterEnabled” on page 1125
- “isEventEnabled” on page 1127
- “listAuditFilters” on page 1128
- “listAuditFiltersByEvent” on page 1129
- “listAuditFiltersByRef” on page 1129
- “modifyAuditFilter” on page 1130

convertFilterRefToString

The convertFilterRefToString command converts a reference ID of a filter to a shortened string value such as AUTHN:SUCCESS.

Target object

None.

Required parameters

-filterRef

Specifies a reference ID for a specific audit filter in the `audit.xml` file. The system defines 4 default audit filters by default. Use the createAuditFilter command to create additional audit filters in your `audit.xml` configuration file. (String, required)

Return value

The command returns the string value of an event type in a shortened format, as the following sample output displays:

```
AUTHN:SUCCESS,AUTHN:INFO,AUTHZ:SUCCESS,AUTHZ:INFO
```

Batch mode example usage

- Using Jython string:

```
AdminTask.convertFilterRefToString('-filterRef AuditSpecification_1184598886859')
```

- Using Jython list:

```
AdminTask.convertFilterRefToString(['-filterRef', 'AuditSpecification_1184598886859'])
```


Interactive mode example usage

- Using Jython:

```
AdminTask.convertFilterRefToString('-interactive')
```

convertFilterStringToRef

The `convertFilterStringToRef` command converts the shortened name of an event type, such as `AUTHN:SUCCESS`, to the reference ID of the audit filter in the `audit.xml` configuration file.

The command accepts one event and outcome pair. The command does not accept multiple event and outcome pairs, such as `AUTHN:SUCCESS AUTHZ:SUCCESS`.

Target object

None.

Required parameters

-filter

Specifies a shortened form of a reference ID for an audit filter, such as `AUTHN:SUCCESS`. The event type must exist in your security auditing system configuration. (String, required)

Return value

The command returns the reference ID for the event type of interest, as the following example displays:

```
AuditSpecification_1173199825608
```

Batch mode example usage

- Using Jython string:

```
AdminTask.convertFilterStringToRef('-filter AUTHN:SUCCESS')
```

- Using Jython list:

```
AdminTask.convertFilterStringToRef(['-filter', 'AUTHN:SUCCESS'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.convertFilterStringToRef('-interactive')
```

createAuditFilter

The `createAuditFilter` command creates and enables a new audit event filter specification entry in the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-name

Specifies a unique name to associate with the audit event filter. (String, required)

-eventType

Specifies a list of one or more auditable events. To specify a list, separate each outcome with a comma (,) character. (String, required)

You can configure the following auditable events in your security auditing system:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities exist.
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out.
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists.
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that could represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates.
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web Services.
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity.

-outcome

Specifies a list of one or multiple event outcomes. For each audit event type, you must specify an outcome. Valid outcomes include SUCCESS, FAILURE, REDIRECT, ERROR, DENIED, WARNING, and INFO. (String, required)

Return value

If the system is successful, the command returns the reference ID for the new audit event filter, as the following sample output displays:

```
AuditSpecification_1184689433421
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditFilter('-name myfilter -eventType  
"SECURITY_MGMT_PROVISIONING, SECURITY_MGMT_POLICY" -outcome SUCCESS')
```

- Using Jython list:

```
AdminTask.createAuditFilter(['-name', 'myfilter', '-eventType',  
"SECURITY_MGMT_PROVISIONING,", 'SECURITY_MGMT_POLICY"', '-outcome', 'SUCCESS'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditFilter('-interactive')
```

deleteAuditFilter

The deleteAuditFilter command deletes the audit event filter specification from the audit.xml file that the system references by an event type and outcome.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-eventType

Specifies the auditable event to delete. (String, required)

The following table displays all valid event types:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities exist.
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out.
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists.
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that could represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates.
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web Services.
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity.

-outcome

Specifies the event outcome. Valid outcomes include SUCCESS, FAILURE, REDIRECT, ERROR, DENIED, WARNING, and INFO. (String, required)

Return value

The command returns a value of true if the system successfully deletes the audit filter from your configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditFilter('-eventType SECURITY_AUTHN -outcome SUCCESS')
```

- Using Jython list:

```
AdminTask.deleteAuditFilter(['-eventType', 'SECURITY_AUTHN', '-outcome', 'SUCCESS'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditFilter('-interactive')
```

deleteAuditFilterByRef

The `deleteAuditFilterByRef` command deletes the audit filter that the system references by the referenced id.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-filterRef

Specifies the reference ID for an audit filter in your security auditing system configuration. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit filter specification from the `audit.xml` file.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditFilterByRef('-filterRef AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.deleteAuditFilterByRef(['-filterRef', 'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditFilterByRef('-interactive')
```

disableAuditFilter

The `disableAuditFilter` command disables the audit filter specification that corresponds to a specific reference id.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-filterRef

Specifies a reference ID for a specific audit filter in the `audit.xml` file. The system defines 4 default audit filters by default. Use the `createAuditFilter` command to create additional audit filters in your `audit.xml` configuration file. (String, required)

Return value

The command returns a value of `true` if the system successfully disables the audit filter.

Batch mode example usage

- Using Jython string:

```
AdminTask.disableAuditFilter('-filterRef', 'AuditSpecification_1184689433421')
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.disableAuditFilter('-interactive')
```

enableAuditFilter

The `enableAuditFilter` command enables the audit filter specification that corresponds to a specific reference id. Use this command to enable a filter that was previously configured and disabled in your security auditing system configuration. To create a new audit filter specification, use the `createAuditFilter` command.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-filterRef

Specifies a reference ID for a specific audit filter in the `audit.xml` file. The system defines 4 default audit filters by default. Use the `createAuditFilter` command to create additional audit filters in your `audit.xml` configuration file. (String, required)

Return value

The command returns a value of `true` if the system successfully enables the audit filter.

Batch mode example usage

- Using Jython string:

```
AdminTask.enableAuditFilter('-filterRef AuditSpecification_1184689433421')
```

- Using Jython list:

```
AdminTask.enableAuditFilter(['-filterRef', 'AuditSpecification_1184689433421'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.enableAuditFilter('-interactive')
```

getAuditFilter

The `getAuditFilter` command retrieves the attributes that the system associates with the audit filter specification of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-reference

Specifies the reference ID for an audit filter in your security auditing system configuration. (String, required)

Return value

The command returns a list of attributes for the audit filter specification of interest, as the following sample output displays:

```
{enabled true}
{name DefaultAuditSpecification_1}
{event SECURITY_AUTHN
SECURITY_AUTHN_MAPPING}
{outcome FAILURE}
[_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditSpecification_1173199825608}
[_Websphere_Config_Data_Type AuditSpecification]}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditFilter('-reference AuditSpecification_1173199825608')
```

- Using Jython list:

```
AdminTask.getAuditFilter(['-reference', 'AuditSpecification_1173199825608'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.getAuditFilter('-interactive')
```

getAuditOutcomes

The `getAuditOutcomes` command retrieves a list of the enabled outcomes for the auditable event type of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventType

Specifies a list of one or more auditable events. To specify a list, separate each outcome with a comma (,) character. (String, required)

You can retrieve the event outcome for any of the following auditable events that might be configured in your security auditing system:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities exist.
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out.

Event name	Description
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists.
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that could represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates.
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web Services.
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity.

Return value

The command returns one or multiple outcomes for the event type of interest, as the following sample output displays:

```
SUCCESS
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditOutcomes('-eventType SECURITY_MGMT_PROVISIONING')
```

- Using Jython list:

```
AdminTask.getAuditOutcomes(['-eventType', 'SECURITY_MGMT_PROVISIONING'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.getAuditOutcomes('-interactive')
```

getSupportedAuditEvents

The `getSupportedAuditEvents` command returns a list of each supported auditable event.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns the following list of possible event outcomes:


```
SECURITY_AUTHN SECURITY_AUTHN_CREDS_MODIFY
SECURITY_AUTHN_DELEGATION SECURITY_AUTHN_MAPPING
SECURITY_AUTHN_TERMINATE SECURITY_AUTHZ
SECURITY_ENCRYPTION SECURITY_MGMT_AUDIT
SECURITY_MGMT_CONFIG SECURITY_MGMT_KEY
SECURITY_MGMT_POLICY SECURITY_MGMT_PROVISIONING
SECURITY_MGMT_REGISTRY SECURITY_MGMT_RESOURCE
SECURITY_RESOURCE_ACCESS SECURITY_RUNTIME
SECURITY_RUNTIME_KEY SECURITY_SIGNING
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getSupportedAuditEvents()
```

- Using Jython list:

```
AdminTask.getSupportedAuditEvents()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getSupportedAuditEvents('-interactive')
```

getSupportedAuditOutcomes

The `getSupportedAuditOutcomes` command retrieves a list of each supported outcome for the auditable event filters.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns the following list of possible event outcomes:

```
SUCCESS
INFO
WARNING
ERROR
DENIED
REDIRECT
FAILURE
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getSupportedAuditOutcomes()
```

- Using Jython list:

```
AdminTask.getSupportedAuditOutcomes()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getSupportedAuditOutcomes('-interactive')
```

isAuditFilterEnabled

The `isAuditFilterEnabled` command determines if the audit filter of interest is enabled in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventType

Specifies a list of one or more auditable events. To specify a list, separate each outcome with a comma (,) character. (String, required)

The following auditable events might be configured in your security auditing system:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities exist.
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out.
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists.
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that could represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates.
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web Services.
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity.

-outcome

Specifies the event outcome. Valid outcomes include SUCCESS, FAILURE, REDIRECT, ERROR, DENIED, WARNING, and INFO. (String, required)

Return value

The command returns a value of true if the event type of interest is enabled in your configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.isAuditFilterEnabled('-eventType SECURITY_MGMT_PROVISIONING  
-outcome SUCCESS')
```

- Using Jython list:

```
AdminTask.isAuditFilterEnabled(['-eventType', 'SECURITY_MGMT_PROVISIONING',  
'-outcome', 'SUCCESS'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.isAuditFilterEnabled('-interactive')
```

isEventEnabled

The `isEventEnabled` command determines if the system enabled at least one audit outcome for the event of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-eventType

Specifies a list of one or more auditable events. To specify a list, separate each outcome with a comma (,) character. (String, required)

The following auditable events are available to configure in your security auditing system:

Event name	Description
SECURITY_AUTHN	Audits all authentication events.
SECURITY_AUTHN_MAPPING	Audits events that record mapping of credentials where two user identities exist.
SECURITY_AUTHN_TERMINATE	Audits authentication termination events such as a timeout, terminated session, or user-initiated logging out.
SECURITY_AUTHZ	Audits events related to authorization checks when the system enforces access control policies.
SECURITY_MGMT_POLICY	Audits events related to security policies, such as the creation of access control lists.
SECURITY_MGMT_PROVISIONING	Audits provisioning events such as the creation an account for a user on a specific machine or adding a user to a group on a specific machine. A given provisioning event might related to one or more SECURITY_MGMT_REGISTRY events.
SECURITY_MGMT_RESOURCE	Audits resource management events such as creation, deletion, and changes to the attributes of a resource. The resource represents an entity with operations that need to be secured. An example of a resource is the Tivoli Access Manager protected object that could represent a file, a Web page, and so on.
SECURITY_RUNTIME_KEY	Audits events related to runtime operations for certificates such as expiration, expiration checks, and invalid certificates.
SECURITY_MGMT_KEY	Audits events related to management operations for certificates such as creating, updating, or exporting a certificate, reading or updating a certificate request, publishing a certificate revocation list, monitoring changes to the keystore, truststore, and so on.
SECURITY_MGMT_AUDIT	Audits events that record operations related to the audit subsystem such as starting audit, stopping audit, turning audit on or off, changing configuration of audit filters or level, archiving audit data, purging audit data, and so on.
SECURITY_RESOURCE_ACCESS	Audits events that record all accesses to a resource. Examples are all accesses to a file, all HTTP requests and responses to a given Web page, and all accesses to a critical database table.
SECURITY_SIGNING	Audits events that record signing such as signing operations used to validate parts of a SOAP Message for Web Services.
SECURITY_ENCRYPTION	Audits events that record encryption information such as encryption for Web services.
SECURITY_AUTHN_DELEGATION	Audits events that record delegation, including Identity Assertion, RunAs, and Low Assertion. Used when the client identity is propagated or when delegation involves the use of a special identity. This event type is also used when switching user identities within a given session.
SECURITY_AUTHN_CREDS_MODIFY	Audits events to modify credentials for a given user identity.

Return value

The command returns a value of `true` if the audit filter of interest has at least one outcome configured in the `audit.xml` file.

Batch mode example usage

- Using Jython string:

```
AdminTask.isEventEnabled('-eventType SECURITY_AUTHN')
```

- Using Jython list:

```
AdminTask.isEventEnabled(['-eventType', 'SECURITY_AUTHN'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isEventEnabled('-interactive')
```

listAuditFilters

The `listAuditFilters` command lists each audit filter and the corresponding attributes that the system defines in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of audit filters and the corresponding attributes, as the following example displays:

```
{(enabled true)
{name DefaultAuditSpecification_1}
{event SECURITY_AUTHN SECURITY_AUTHN_MAPPING}
{outcome FAILURE}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1173199825608}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1173199825608}}
{(enabled true)
{name DefaultAuditSpecification_2}
{event {}}
{outcome FAILURE}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1173199825609}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1173199825609}}
{(enabled true)
{name DefaultAuditSpecification_3}
{event SECURITY_RESOURCE_ACCESS}
{outcome FAILURE}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1173199825610}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1173199825610}}
{(enabled true)
{name DefaultAuditSpecification_4}
{event SECURITY_AUTHN_TERMINATE}
{outcome FAILURE}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1173199825611}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1173199825611}}
{(enabled true)
{name myfilter}
{event SECURITY_AUTHZ}
{outcome REDIRECT}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184365235250}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1184365235250}}
{(enabled true)
{name myfilter1}
{event SECURITY_AUTHZ SECURITY_RESOURCE_ACCESS}
{outcome REDIRECT INFO}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184365353218}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1184365353218}}
{(enabled true)
{name myfilter}
{event SECURITY_AUTHN SECURITY_AUTHZ}
{outcome SUCCESS INFO}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184598886859}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1184598886859}}
{(enabled false)
{name myfilter}
{event SECURITY_MGMT_PROVISIONING SECURITY_MGMT_POLICY}}
```

```
{outcome SUCCESS}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#AuditSpecification_1184689433421}
{_Websphere_Config_Data_Type AuditSpecification}
{filterRef AuditSpecification_1184689433421}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditFilters()
```

- Using Jython list:

```
AdminTask.listAuditFilters()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditFilters('-interactive')
```

listAuditFiltersByEvent

The `listAuditFiltersByEvent` command retrieves a list of events and event outcomes for each audit filter that is configured in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of events and event outcomes for the audit filters of interest, as the following sample output displays:

```
{AuditSpecification_1173199825608 SECURITY_AUTHN:FAILURE}{AuditSpecification_1173199825610 SECURITY_RESOURCE_ACCESS:FAILURE}{AuditSpecification_1173199825611 SECURITY_AUTHN_TERMINATE:FAILURE}{AuditSpecification_1184365235250 SECURITY_AUTHZ:REDIRECT}{AuditSpecification_1184365353218 SECURITY_AUTHZ:REDIRECT;SECURITY_AUTHZ:INFO}{AuditSpecification_1184365353218 SECURITY_RESOURCE_ACCESS:REDIRECT;SECURITY_RESOURCE_ACCESS:INFO}{AuditSpecification_1184598886859 SECURITY_AUTHN:SUCCESS;SECURITY_AUTHN:INFO}{AuditSpecification_1184598886859 SECURITY_AUTHZ:SUCCESS;SECURITY_AUTHZ:INFO}{AuditSpecification_1184689433421 SECURITY_MGMT_PROVISIONING:SUCCESS}{AuditSpecification_1184689433421 SECURITY_MGMT_POLICY:SUCCESS}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditFiltersByEvent()
```

- Using Jython list:

```
AdminTask.listAuditFiltersByEvent()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditFiltersByEvent('-interactive')
```

listAuditFiltersByRef

The `listAuditFiltersByRef` command lists all reference ids that correspond to the audit filters that are defined in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of each reference that exists in the audit.xml configuration file, as the following sample output displays:

```
AuditSpecification_1173199825608 AuditSpecification_1173199825609
AuditSpecification_1173199825610 AuditSpecification_1173199825611
AuditSpecification_1184365235250 AuditSpecification_1184365353218
AuditSpecification_1184598886859 AuditSpecification_1184689433421
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditFiltersByRef()
```

- Using Jython list:

```
AdminTask.listAuditFiltersByRef()
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.listAuditFiltersByRef('-interactive')
```

modifyAuditFilter

The modifyAuditFilter command modifies the audit filter specification in the audit.xml configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-filterRef

Specifies the reference ID for the audit filter to modify in your security auditing system configuration. (String, required)

Optional parameters

-name

Specifies a unique name to associate with the audit event filter. (String, optional)

-eventType

Specifies a comma-separated list of one or more event types. (String, optional)

-outcome

Specifies a comma-separated list of one or multiple event outcomes. For each audit event type, you must specify an outcome. Valid outcomes include SUCCESS, FAILURE, REDIRECT, ERROR, DENIED, WARNING, and INFO. (String, optional)

-enableFilter

Specifies whether to enable the filter. Specify true to enable the filter, or false to disable the filter. (Boolean, optional).

Return value

The command returns a value of true if the system successfully updates the audit filter.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditFilter('-filterRef AuditSpecification_1173199825608 -name myname -eventType SECURITY_AUTHN -outcome SUCCESS -enableFilter true')
```

- Using Jython list:

```
AdminTask.modifyAuditFilter(['-filterRef', 'AuditSpecification_1173199825608', '-name', 'myname', '-eventType', 'SECURITY_AUTHN', '-outcome', 'SUCCESS', '-enableFilter', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditFilter('-interactive')
```

AuditNotificationCommands command group for the AdminTask object

You can use the Jython scripting language to manage the security auditing system with the wsadmin tool. Use the commands and parameters in the AuditNotificationCommands group to configure and manage audit notifications and audit notification monitors.

Use the following commands to configure your security auditing system notifications:

- “createAuditNotification”
- “createAuditNotificationMonitor” on page 1132
- “deleteAuditNotification” on page 1133
- “deleteAuditNotificationMonitorByName” on page 1133
- “deleteAuditNotificationMonitorByRef” on page 1134
- “getAuditNotification” on page 1135
- “getAuditNotificationMonitor” on page 1135
- “getEmailList” on page 1136
- “getSendEmail” on page 1136
- “getAuditNotificationRef” on page 1137
- “getAuditNotificationName” on page 1137
- “isSendEmailEnabled” on page 1138
- “isAuditNotificationEnabled” on page 1138
- “listAuditNotifications” on page 1139
- “listAuditNotificationMonitors” on page 1139
- “modifyAuditNotification” on page 1140
- “modifyAuditNotificationMonitor” on page 1141
- “setEmailList” on page 1141
- “setSendEmail” on page 1142

createAuditNotification

The createAuditNotification command creates an audit notification object in the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-notificationName

Specifies the unique name to assign the audit notification object. (String, required)

-logToSystemOut

Specifies whether the system logs notifications to the `SystemOut.log` file. (Boolean, required)

-sendEmail

Specifies whether to email security auditing subsystem failure notifications. (Boolean, required)

Optional parameters

-emailList

Specifies the email list to send security auditing subsystem failure notifications. (String, optional)

-emailFormat

Specifies the email format. Specify HTML for HTML format or TEXT for text format. (String, optional)

Return value

The command returns the shortened reference ID of the new audit notification object, as the following sample output displays:

```
WSNotification_1184690835390
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditNotification(['-notificationName mynotification -logToSystemOut true -sendEmail true -emailList admin@mycompany.com(smtp-server.mycompany.com) -emailFormat HTML'])
```

- Using Jython list:

```
AdminTask.createAuditNotification(['-notificationName', 'mynotification', '-logToSystemOut', 'true', '-sendEmail', 'true', '-emailList', 'admin@mycompany.com(smtp-server.mycompany.com)', '-emailFormat', 'HTML'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditNotification('-interactive')
```

createAuditNotificationMonitor

The `createAuditNotificationMonitor` command creates an audit notification monitor object for the security auditing system. This object monitors the security auditing subsystem for possible failure.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-monitorName

Specifies the unique name of the audit notification monitoring object. (String, required)

-notificationRef

Specifies the reference ID of the audit notification object. (String, required)

-enable

Specifies whether to enable the audit notification monitor. (Boolean, required)

Return value

The command returns the shortened form of the reference ID for the audit notification monitor, as the following sample output displays:

```
AuditNotificationMonitor_1184695615171
```

Batch mode example usage

- Using Jython string:

```
AdminTask.createAuditNotificationMonitor('-monitorName mymonitor -notificationRef  
WSNotification_1184690835390 -enable true')
```

- Using Jython list:

```
AdminTask.createAuditNotificationMonitor(['-monitorName', 'mymonitor', '-notificationRef',  
'WSNotification_1184690835390', '-enable', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.createAuditNotificationMonitor('-interactive')
```

deleteAuditNotification

The `deleteAuditNotification` command deletes an audit notification object from the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-notificationRef

Specifies the reference ID of the audit notification object to delete. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit notification object from the `audit.xml` configuration file.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditNotification('-notificationRef WSNotification_1184690835390')
```

- Using Jython list:

```
AdminTask.deleteAuditNotification(['-notificationRef', 'WSNotification_1184690835390'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditNotification('-interactive')
```

deleteAuditNotificationMonitorByName

The `deleteAuditNotificationMonitorByName` command deletes the audit notification monitor that the user specifies with the unique name.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-monitorName

Specifies the unique name of the audit notification monitor to delete. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit notification monitor from the configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditNotificationMonitor('-monitorName mymonitor')
```

- Using Jython list:

```
AdminTask.deleteAuditNotificationMonitor(['-monitorName', 'mymonitor'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditNotificationMonitor('-interactive')
```

deleteAuditNotificationMonitorByRef

The `deleteAuditNotificationMonitorByRef` command deletes the audit notification monitor that the user specifies with the reference ID.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-monitorRef

Specifies the reference ID of the audit notification monitor object to delete. (String, required)

Return value

The command returns a value of `true` if the system successfully deletes the audit notification monitor of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.deleteAuditNotificationMonitor('-monitorRef  
AuditNotificationMonitor_1184695615171')
```

- Using Jython list:

```
AdminTask.deleteAuditNotificationMonitor(['-monitorRef',  
'AuditNotificationMonitor_1184695615171'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.deleteAuditNotificationMonitor('-interactive')
```

getAuditNotification

The `getAuditNotification` command retrieves the attributes for an audit notification object of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-notificationRef

Specifies the reference ID of the audit notification object of interest. (String, required)

Return value

The command returns a list of attributes for the specific audit notification object, as the following sample output displays:

```
{name mynotification}
{sslConfig {}}
{logToSystemOut true}
{_Websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#WSNotification_1184690835390}
{emailList sweetshadow@us.ibm.com(smtp-server.us.ibm.com)}
{sendEmail true}
{_Websphere_Config_Data_Type WSNotification}
{properties {}}
{emailFormat HTML}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditNotification('-notificationRef WSNotification_1184690835390')
```

- Using Jython list:

```
AdminTask.getAuditNotification(['-notificationRef', 'WSNotification_1184690835390'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditNotification('-interactive')
```

getAuditNotificationMonitor

The `getAuditNotificationMonitor` command retrieves the attributes that the system associates with the audit notification monitor of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-monitorRef

Specifies the reference ID of the audit notification monitor of interest. (String, required)

Return value

The command returns a list of attributes for the audit notification monitor of interest, as the following sample output displays:

```
{{name mymonitor}
{enabled true}
{Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditNotificationMonitor_1184695615171}
{Websphere_Config_Data_Type AuditNotificationMonitor}
{wsNotification mynotification(cells/Node04Cell|audit.xml#WSNotification_1184690835390)}}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditNotificationMonitor('-monitorRef',
AuditNotificationMonitor_1184695615171')
```

- Using Jython list:

```
AdminTask.getAuditNotificationMonitor(['-monitorRef',
'AuditNotificationMonitor_1184695615171'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditNotificationMonitor('-interactive')
```

getEmailList

The getEmailList command retrieves the email distribution list for the audit notification object. If the notification monitor is not configured, the audit notification object is not active and the command returns a null value.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns this email list for the active audit notification object, as the following sample output displays:

```
admin@mycompany.com(smtp-server.mycompany.com)
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getEmailList()
```

- Using Jython list:

```
AdminTask.getEmailList()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEmailList('-interactive')
```

getSendEmail

The getSendEmail command displays whether or not the audit notification object sends an email if the audit subsystem fails. If the notification monitor is not configured, the audit notification object is not active and the command returns a null value.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system is configured to send an email to the distribution list.

Batch mode example usage

- Using Jython string:

```
AdminTask.getSendEmail()
```

- Using Jython list:

```
AdminTask.getSendEmail()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getSendEmail('-interactive')
```

getAuditNotificationRef

The `getAuditNotificationRef` command retrieves the reference ID for the active audit notification object. If the notification monitor is not configured, the audit notification object is not active and the command returns a null value.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns the reference ID of the audit notification object if it is active, as the following sample output displays:

```
WSNotification_1184690835390
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditNotificationRef()
```

- Using Jython list:

```
AdminTask.getAuditNotificationRef()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditNotificationRef('-interactive')
```

getAuditNotificationName

The `getAuditNotificationName` command retrieves the unique name for the active audit notification object. If the notification monitor is not configured, the audit notification object is not active and the command returns a null value.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns the unique name of the audit notification object, as the following sample output displays:

```
mynotification
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditNotificationName()
```

- Using Jython list:

```
AdminTask.getAuditNotificationName()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditNotificationName('-interactive')
```

isSendEmailEnabled

The `isSendEmailEnabled` command determines if the system is configured to send an email if the security auditing subsystem fails.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if email notification is enabled.

Batch mode example usage

- Using Jython string:

```
AdminTask.isSendEmailEnabled()
```

- Using Jython list:

```
AdminTask.isSendEmailEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isSendEmailEnabled('-interactive')
```

isAuditNotificationEnabled

The `isAuditNotificationEnabled` command determines whether the security auditing system notifications are enabled.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if security auditing system notifications are enabled.

Batch mode example usage

- Using Jython string:

```
AdminTask.isAuditNotificationEnabled()
```

- Using Jython list:

```
AdminTask.isAuditNotificationEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isAuditNotificationEnabled()
```

listAuditNotifications

The `listAuditNotifications` command retrieves the attributes for each audit notification object that is configured in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes for each configured audit notification object, as the following sample output displays:

```
{{name mynotification}
{sslConfig {}}
{logToSystemOut true}
{websphere_Config_Data_Id cells/CHEYENNENode04Cell|audit.xml#WSNotification_1184690835390}
{emailList sweetshadow@us.ibm.com(smtp-server.us.ibm.com)}
{sendEmail true}
{notificationRef WSNotification_1184690835390}
{websphere_Config_Data_Type WSNotification}
{properties {}}
{emailFormat HTML}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditNotifications()
```

- Using Jython list:

```
AdminTask.listAuditNotifications()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditNotifications('-interactive')
```

listAuditNotificationMonitors

The `listAuditNotificationMonitors` command lists the attributes for the audit notification monitor that is configured in the `audit.xml` file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes for the audit notification monitor, as the following sample output displays:

```
{{name mymonitor}
{enabled true}
{_Websphere_Config_Data_Id cells/Node04Cell|audit.xml#AuditNotificationMonitor_1184695615171}
{_Websphere_Config_Data_Type AuditNotificationMonitor}
{monitorRef AuditNotificationMonitor_1184695615171}
{wsNotification mynotification(cells/Node04Cell|audit.xml#WSNotification_1184690835390)}
{notificationRef WSNotification_1184690835390}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.listAuditNotificationMonitors()
```

- Using Jython list:

```
AdminTask.listAuditNotificationMonitors()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.listAuditNotificationMonitors('-interactive'b)
```

modifyAuditNotification

The `modifyAuditNotification` command edits the audit notification object in the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-notificationRef

Specifies the reference ID of the audit notification object to edit. (String, required)

Optional parameters

-logToSystemOut

Specifies whether to log notifications to the `SystemOut.log` file. (Boolean, optional)

-sendEmail

Specifies whether to email notifications. (Boolean, optional)

-emailList

Specifies the email address of distribution list where the system sends email notifications. (String, optional)

-emailFormat

Specifies the email format. Specify HTML for HTML format or TEXT for text format. (String, optional)

Return value

The command returns a value of `true` if the system successfully updates the security auditing system configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditNotification('-notificationRef WSNotification_1184690835390
-logToSystemOut false -sendEmail true -emailList admin@mycompany.com(smtp-server.mycompany.com)
-emailFormat TEXT')
```

- Using Jython list:

```
AdminTask.modifyAuditNotification(['-notificationRef', 'WSNotification_1184690835390',
'-logToSystemOut', 'false', '-sendEmail', 'true', '-emailList',
'admin@mycompany.com(smtp-server.mycompany.com)', '-emailFormat', 'TEXT'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditNotification('-interactive')
```

modifyAuditNotificationMonitor

The `modifyAuditNotificationMonitor` command edits the audit notification monitor configuration for the security auditing system.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-monitorRef

Specifies the reference ID of the audit notification monitor of interest. (String, required)

Optional parameters

-notificationRef

Specifies the reference ID of the audit notification object. (String, optional)

-enable

Specifies whether to enable the audit notification monitor. (Boolean, optional)

Return value

The command returns a value of `true` if the system successfully updates the audit notification monitor configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditNotificationMonitor('-monitorRef AuditNotificationMonitor_1184695615171
-notificationRef WSNotification_1184690835390 -enable true')
```

- Using Jython list:

```
AdminTask.modifyAuditNotificationMonitor(['-monitorRef', 'AuditNotificationMonitor_1184695615171',
'-notificationRef', 'WSNotification_1184690835390', '-enable', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditNotificationMonitor('-interactive')
```

setEmailList

The `setEmailList` command specifies the distribution list to send email notifications to if the security auditing subsystem fails.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-emailList

Specifies the email address or email distribution list to send audit notifications to. (String, required)

Return value

The command returns a value of `true` if the system successfully sets the email notification list for the notification object.

Batch mode example usage

- Using Jython string:

```
AdminTask.setEmailList(['-emailList admin@mycompany.com(smtp-server.mycompany.com)'])
```

- Using Jython list:

```
AdminTask.setEmailList(['-emailList', 'admin@mycompany.com(smtp-server.mycompany.com)'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.setEmailList('-interactive')
```

setSendEmail

The `setSendEmail` command enables or disables email notifications for the security auditing system.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-enable

Specifies whether to enable the system to send audit notifications by email. (Boolean, required)

Return value

The command returns a value of `true` if the system successfully modifies the configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.setSendEmail('-enable true')
```

- Using Jython list:

```
AdminTask.setSendEmail(['-enable', 'true'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setSendEmail('-interactive')
```

AuditPolicyCommands command group for the AdminTask object

You can use the Jython scripting language to manage the security auditing system with the `wsadmin` tool. Use the commands and parameters in the `AuditPolicyCommands` group to enable and configure the security auditing system.

Use the following commands to configure, query, and manage the security auditing system:

- “disableAudit”
- “disableVerboseAudit”
- “enableAudit” on page 1144
- “enableVerboseAudit” on page 1144
- “getAuditPolicy” on page 1145
- “getAuditSystemFailureAction” on page 1146
- “getAuditorId” on page 1146
- “isAuditEnabled” on page 1147
- “isVerboseAuditEnabled” on page 1147
- “mapAuditGroupIDsOfAuthorizationGroup” on page 1148
- “modifyAuditPolicy” on page 1148
- “setAuditSystemFailureAction” on page 1149
- “resetAuditSystemFailureAction” on page 1150
- “setAuditorId” on page 1150
- “setAuditorPwd” on page 1151

disableAudit

The disableAudit command disables security auditing in the `audit.xml` configuration file.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully disables security auditing.

Batch mode example usage

- Using Jython string:

```
AdminTask.disableAudit()
```

- Using Jython list:

```
AdminTask.disableAudit()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.disableAudit('-interactive')
```

disableVerboseAudit

The disableVerboseAudit command disables the verbose capture of audit data for the security auditing system.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully disables the verbose capture of audit data.

Batch mode example usage

- Using Jython string:

```
AdminTask.disableVerboseAudit()
```

- Using Jython list:

```
AdminTask.disableVerboseAudit()
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.disableVerboseAudit('-interactive')
```

enableAudit

The `enableAudit` command enables security auditing in the `audit.xml` configuration file. By default, security auditing is disabled.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully enables security auditing.

Batch mode example usage

- Using Jython string:

```
AdminTask.enableAudit()
```

- Using Jython list:

```
AdminTask.enableAudit()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.enableAudit('-interactive')
```

enableVerboseAudit

The `enableVerboseAudit` command sets the security auditing system to perform verbose capture of audit data.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully sets the security auditing system to perform verbose capture of audit data.

Batch mode example usage

- Using Jython string:

```
AdminTask.enableVerboseAudit()
```

- Using Jython list:

```
AdminTask.enableVerboseAudit()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.enableVerboseAudit('-interactive')
```

getAuditPolicy

The `getAuditPolicy` command retrieves each attribute that is associated with the audit policy in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a list of attributes for the security auditing system, as the following sample output displays:

```
{auditEventFactories {{name auditEventFactoryImpl_1
{properties {}
{className com.ibm.ws.security.audit.AuditEventFactoryImpl}
{auditServiceProvider auditServiceProviderImpl_1(cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608)}
{auditSpecifications DefaultAuditSpecification_1(cells/Node04Cell|audit.xml#AuditSpecification_1173199825608)
DefaultAuditSpecification_2(cells/Node04Cell|audit.xml#AuditSpecification_1173199825609)
DefaultAuditSpecification_3(cells/Node04Cell|audit.xml#AuditSpecification_1173199825610)
DefaultAuditSpecification_4(cells/Node04Cell|audit.xml#AuditSpecification_1173199825611)}
{websphere_config_data_id cells/Node04Cell|audit.xml#AuditEventFactory_1173199825608}
{websphere_config_data_type AuditEventFactory}}}}
{websphere_config_data_id cells/Node04Cell|audit.xml#AuditPolicy_1173199825608}
{auditServiceProviders {{auditSpecifications
DefaultAuditSpecification_1(cells/Node04Cell|audit.xml#AuditSpecification_1173199825608)
DefaultAuditSpecification_2(cells/Node04Cell|audit.xml#AuditSpecification_1173199825609)
DefaultAuditSpecification_3(cells/Node04Cell|audit.xml#AuditSpecification_1173199825610)
DefaultAuditSpecification_4(cells/Node04Cell|audit.xml#AuditSpecification_1173199825611)}
{name auditServiceProviderImpl_1
{websphere_config_data_id cells/Node04Cell|audit.xml#AuditServiceProvider_1173199825608}
{maxFileSize 1}
{websphere_config_data_type AuditServiceProvider}
{fileLocation ${PROFILE_ROOT}/logs/server1}
{className com.ibm.ws.security.audit.BinaryEmitterImpl}
{properties {}
{eventFormatterClass {}
{maxLogs 100}}}}
{securityXmlSignerCertAlias auditSignCert}
{properties {}
{securityXmlSignerScopeName (cell):Node04Cell:(node):Node04}
{auditorPwd SweetShadowsPwd}
{websphere_config_data_type AuditPolicy}
{securityXmlSignerKeyStoreName NodeDefaultSignersStore}
{verbose false}
{auditPolicy WARN}
{encrypt false}
{managementScope {}
{encryptionCert {}
{batching false}
{auditorId SweetShadow}
{auditEnabled false}
{sign true}}
```

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditPolicy()
```

- Using Jython list:

AdminTask.getAuditPolicy()

Interactive mode example usage

- Using Jython:

AdminTask.getAuditPolicy('-interactive')

getAuditSystemFailureAction

The `getAuditSystemFailureAction` command displays the action that the application server takes if a failure occurs in the security auditing subsystem.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a string that describes the action that the application server takes if the security auditing subsystem fails. Possible values are `WARN`, `NOWARN`, or `FATAL`. The following table describes the behavior associated with each action that the application server takes if the security auditing subsystem fails:

WARN	Specifies that the application server should notify the auditor, stop security auditing, and continue to run the application server process.
NOWARN	Specifies that the application server should not notify the auditor, but should stop security auditing and continue to run the application server process
FATAL	Specifies that the application server should notify the auditor, stop security auditing, and stop the application server process.

Batch mode example usage

- Using Jython string:

AdminTask.getAuditSystemFailureAction()

- Using Jython list:

AdminTask.getAuditSystemFailureAction()

Interactive mode example usage

- Using Jython:

AdminTask.getAuditSystemFailureAction('-interactive')

getAuditorId

The `getAuditorId` command retrieves the name of the user who is assigned as the auditor.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns the name of the user who is assigned as the auditor.

Batch mode example usage

- Using Jython string:

```
AdminTask.getAuditorId()
```

- Using Jython list:

```
AdminTask.getAuditorId()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getAuditorId('-interactive')
```

isAuditEnabled

The `isAuditEnabled` command determines whether the security auditing is enabled in your configuration. By default, auditing is not enabled in the `audit.xml` configuration file.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if security auditing is enabled in your environment. If the command returns a value of `false`, security auditing is disabled.

Batch mode example usage

- Using Jython string:

```
AdminTask.isAuditEnabled()
```

- Using Jython list:

```
AdminTask.isAuditEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isAuditEnabled('-interactive')
```

isVerboseAuditEnabled

The `isVerboseAuditEnabled` command determines whether or not the security auditing system verbosely captures audit data.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the security auditing system is configured to verbosely capture audit data.

Batch mode example usage

- Using Jython string:

```
AdminTask.isVerboseAuditEnabled()
```

- Using Jython list:

```
AdminTask.isVerboseAuditEnabled()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.isVerboseAuditEnabled('-interactive')
```

mapAuditGroupIDsOfAuthorizationGroup

The mapAuditGroupIDsOfAuthorizationGroup command maps the special subjects to users in the registry.

The user must have the monitor administrative role to run this command.

Target object

None.

Return value

The command does not return output.

Batch mode example usage

- Using Jython string:

```
AdminTask.mapAuditGroupIDsOfAuthorizationGroup()
```

- Using Jython list:

```
AdminTask.mapAuditGroupIDsOfAuthorizationGroup()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.mapAuditGroupIDsOfAuthorizationGroup('-interactive')
```

modifyAuditPolicy

The modifyAuditPolicy command modifies the audit policy attributes in the audit.xml configuration file. You can use this command to modify one or multiple attributes.

The user must have the auditor administrative role to run this command.

Target object

None.

Optional parameters

-auditEnabled

Specifies whether security auditing is enabled in your configuration. (Boolean, optional)

-auditPolicy

Specifies the action that the application server takes if the security auditing subsystem fails. (String, optional)

The following table describes the valid values for this parameter:

WARN	Specifies that the application server should notify the auditor, stop security auditing, and continue to run the application server process.
NOWARN	Specifies that the application server should not notify the auditor, but should stop security auditing and continue to run the application server process

FATAL	Specifies that the application server should notify the auditor, stop security auditing, and stop the application server process.
-------	---

-auditorId

Specifies the name of the user that the system assigns as the auditor. (String, optional)

-auditorPwd

Specifies the password for the auditor id. (String, optional)

-sign

Specifies whether to sign audit records. Use the AuditSigningCommands command group to configure signing settings. (Boolean, optional)

-encrypt

Specifies whether to encrypt audit records. Use the AuditEncryptionCommands command group to configure encryption settings. (Boolean, optional)

-verbose

Specifies whether to capture verbose audit data. (Boolean, optional)

-encryptionCert

Specifies the reference ID of the certificate to use for encryption. Specify this parameter if you set the -encrypt parameter to true. (String, optional)

Return value

The command returns a value of true if the system successfully updates the security auditing system policy.

Batch mode example usage

- Using Jython string:

```
AdminTask.modifyAuditPolicy('-auditEnabled true -auditPolicy NOWARN -auditorId testuser -auditorPwd testuserpwd -sign false -encrypt false -verbose false')
```

- Using Jython list:

```
AdminTask.modifyAuditPolicy(['-auditEnabled', 'true', '-auditPolicy', 'NOWARN', '-auditorId', 'testuser', '-auditorPwd', 'testuserpwd', '-sign', 'false', '-encrypt', 'false', '-verbose', 'false'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.modifyAuditPolicy('-interactive')
```

setAuditSystemFailureAction

The setAuditSystemFailureAction command sets the action that the application server takes if the security auditing subsystem fails.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-action

Specifies the action to take if the security auditing subsystem fails. (String, required)

The following table describes the valid values for this parameter:

WARN	Specifies that the application server should notify the auditor, stop security auditing, and continue to run the application server process.
NOWARN	Specifies that the application server should not notify the auditor, but should stop security auditing and continue to run the application server process
FATAL	Specifies that the application server should notify the auditor, stop security auditing, and stop the application server process.

Return value

The command returns a value of `true` if the system successfully updates the security auditing system policy.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAuditSystemFailureAction('-action NOWARN')
```

- Using Jython list:

```
AdminTask.setAuditSystemFailureAction(['-action', 'NOWARN'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setAuditSystemFailureAction('-interactive')
```

resetAuditSystemFailureAction

The `resetAuditSystemFailureAction` command sets the action that the application server takes if the security auditing system fails to the `NOWARN` setting.

The user must have the auditor administrative role to run this command.

Target object

None.

Return value

The command returns a value of `true` if the system successfully updates your configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.resetAuditSystemFailureAction()
```

- Using Jython list:

```
AdminTask.resetAuditSystemFailureAction()
```

Interactive mode example usage

- Using Jython:

```
AdminTask.resetAuditSystemFailureAction('-interactive')
```

setAuditorId

The `setAuditorId` command sets the name of the user to assign as the auditor.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-name

Specifies the name of the user to assign as the auditor. (String, required)

Return value

The command returns a value of `true` if the system successfully updates your configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAuditorId('-name myAdmin')
```

- Using Jython list:

```
AdminTask.setAuditorId(['-name', 'myAdmin'])
```

Interactive mode example usage

- Using Jython string:

```
AdminTask.setAuditorId('-interactive')
```

setAuditorPwd

The `setAuditorPwd` command sets the password for the auditor.

The user must have the auditor administrative role to run this command.

Target object

None.

Required parameters

-password

Specifies the password for the user assigned as the auditor. (String, required)

Return value

The command returns a value of `true` if the system successfully updates your configuration.

Batch mode example usage

- Using Jython string:

```
AdminTask.setAuditorPwd('-password myAdminPassword')
```

- Using Jython list:

```
AdminTask.setAuditorPwd(['-password', 'myAdminPassword'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.setAuditorPwd('-interactive')
```

AuditEventFormatterCommands command group for the AdminTask object

You can use the Jython scripting language to manage the security auditing system with the `wsadmin` tool. Use the commands and parameters in the `AuditEventFormatterCommands` group to manage the event formatter for the audit service provider.

Use the following commands to display information about the audit event formatter:

- “getEventFormatterClass”

getEventFormatterClass

The `getEventFormatterClass` command retrieves the class name of the event formatter specified by the reference ID of the binary file audit service provider of interest.

The user must have the monitor administrative role to run this command.

Target object

None.

Required parameters

-emitterRef

Specifies the reference ID of the audit service provider implementation of interest. (String, required)

Return value

The command returns a null value if the event formatter class is not defined for the audit service provider implementation of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.getEventFormatterClass('-emitterRef AuditServiceProvider_1173199825608')
```

- Using Jython list:

```
AdminTask.getEventFormatterClass(['-emitterRef', 'AuditServiceProvider_1173199825608'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.getEventFormatterClass('-interactive')
```

AuditReaderCommands command group for the AdminTask object

You can use the Jython scripting language to manage the security auditing system with the `wsadmin` tool. Use the commands and parameters in the `AuditReaderCommands` group to display audit record information from the binary audit log.

Use the following commands to query the binary audit log:

- “binaryAuditLogReader”
- “showAuditLogEncryptionInfo” on page 1154

binaryAuditLogReader

The `binaryAuditLogReader` command reads the default binary audit log and generates an HTML report based on the parameters you provide. You must use the auditor security role to use this command.

Target object

None.

Required parameters

-fileName

Specifies the fully qualified file name for the binary audit log. (String, required)

-outputLocation

Specifies the location of the HTML report that the command generates. (String, required)

Optional parameters

-reportMode

Specifies the type of report to generate. Valid values include `basic`, `complete`, or `custom`. The `basic` report provides the following configuration information:

- `creationTime`
- `action`
- `progName`
- `registryType`
- `domain`
- `realm`
- `remoteAddr`
- `remotePort`
- `remoteHost`
- `resourceName`
- `resourceType`
- `resourceUniqueld`

The `complete` report provides the data included by the default report type and each additional datapoint of interest. The `custom` report allows you to specify only the datapoints you choose to see generated in the report. The default value is `basic`. (String, optional)

-eventFilter

Specifies the audit types to read and report. Specify one or more audit event types. If you specify more than one value for the `eventFilter` parameter, separate each audit event type with a colon character (:). (String, optional)

-outcomeFilter

Specifies the audit event outcomes to read and report. Specify one or more audit event outcomes. If you specify more than one value for the `outcomeFilter` parameter, separate each audit event outcome with a colon character (:). (String, optional)

-sequenceFilter

Specifies a list of beginning and ending sequence numbers. Use the `a:b` syntax, where `a`, the starting sequence number where the HTML report begins, and is less than or equal to `b`, the sequence number where the HTML report ends. A single sequence may also be specified, such as `-sequenceFilter 10`, to only generate a report for the tenth record. (String, optional)

-timestampFilter

Specifies the time stamp range of records to read and report. Use the `a:b` syntax, where `a` and `b` are strings in the format `java.text.SimpleDateFormat("MMddhhmmyyy")`. You can also specify a single timestamp. (String, optional)

-keyStorePassword

Specifies password to open the keystore. (String, optional)

-dataPoints

Specifies a list of specific audit data to use to generate the report. Use this option only when you set the `reportMode` parameter as `custom`. If you specify multiple data points, separate each data point with a colon character (:). (String, optional)

Return value

The command returns the HTML report based on the values specified for each parameter to the location specified by the `outputLocation` parameter.

Batch mode example usage

- Using Jython string:

Windows

```
AdminTask.binaryAuditLogReader(['-fileName myFileName -reportMode basic  
-keyStorePassword password123 -outputLocation C:\binaryLogs'])
```

AIX

HP-UX

Solaris

Linux

```
AdminTask.binaryAuditLogReader(['-fileName myFileName -reportMode basic  
-keyStorePassword password123 -outputLocation /binaryLogs'])
```

- Using Jython list:

Windows

```
AdminTask.binaryAuditLogReader(['-fileName', 'myFileName', '-reportMode', 'basic',  
'-keyStorePassword', 'password123', '-outputLocation', 'C:\binaryLogs'])
```

AIX

HP-UX

Solaris

Linux

```
AdminTask.binaryAuditLogReader(['-fileName', 'myFileName', '-reportMode', 'basic',  
'-keyStorePassword', 'password123', '-outputLocation', '/binaryLogs'])
```

Interactive mode example usage

- Using Jython:

```
AdminTask.binaryAuditLogReader('-interactive')
```

showAuditLogEncryptionInfo

The `showAuditLogEncryptionInfo` command displays information about the keystore that the auditing system uses to encrypt audit records. Use this information as a hint of the keystore password in order to decrypt encrypted audit logs in the binary audit log.

Target object

None.

Required parameters

-fileName

Specifies the fully qualified path of the binary audit log. (String, required)

Return value

The command returns the certificate alias and the fully qualified path to the keystore of interest.

Batch mode example usage

- Using Jython string:

```
AdminTask.showAuditLogEncryptionInfo('-fileName myFileName')
```

- Using Jython list:

```
AdminTask.showAuditLogEncryptionInfo(['-fileName', 'myFileName'])
```

Interactive mode example usage

- Using Jython:

SSLMigrationCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to migrate key store configurations. Use the commands in the SSLMigrationCommands group to convert self-signed certificates to chained personal certificates and to enable writable key rings.

The SSLMigrationCommands command group for the AdminTask object includes the following commands:

- “convertSelfSignedCertificatesToChained command”
- “enableWritableKeyrings command” on page 1156
- “convertSSLConfig command” on page 1157

convertSelfSignedCertificatesToChained command

The convertSelfSignedCertificatesToChained command converts specific self-signed certificates to chained personal certificates.

Note: Chained certificates are the default certificate type in Websphere Application Server Version 7.0. The convertSelfSignedCertificatesToChained command takes information from the self-signed certificate—such as issued-to DN, size, and life span—and creates a chained certificate with the same information. The new chained certificate replaces the self-signed certificate. Signer certificates from the self-signed certificate that are distributed across the security configuration are replaced with the signer certificates from the root certificate used to sign the chained certificate.

Syntax

The command has the following syntax:

```
wsadmin>$AdminTask convertSelfSignedCertificatesToChained
    [-certificateReplacementOption ALL_CERTIFICATES | DEFAULT_CERTIFICATES | KEYSTORE_CERTIFICATES]
    [-keyStoreName keystore_name]
    [-keyStoreScope keystore_scope]
    [-rootCertificateAlias alias_name]
```

Required parameters

certificateReplacementOption

Specifies the convert self-signed certificates replacement options. (String, required)

Specify the value for the parameter as one of the following options:

ALL_CERTIFICATES

This option looks for all self-signed certificates in all keystores with in the specified scope.

The scope can be provided in the -keyStoreScope parameter. If no scope is provided using the -keyStoreScope parameter, all scopes are visited.

DEFAULT_CERTIFICATES

This option looks for self-signed certificates in the default CellDefaultKeyStore and NodeDefaultKeyStore keystores within the specified scope.

The scope can be provided with the -keyStoreScope parameter. If no scope is provided using the -keyStoreScope parameter, all scopes are visited.

KEYSTORE_CERTIFICATES

This option replaces only those self-signed certificates in the keystore that are specified by the -keyStoreName parameter.

If no scope is provided using the `-keyStoreScope` parameter, the default scope is used.

Optional parameters

keyStoreName

Specifies the name of a keystore in which to look for self-signed certificates to convert. Use this parameter with the `KEYSTORE_CERTIFICATES` option on the `certificateReplacementOption` parameter. (String, optional)

keyStoreScope

Specifies the name of the scope in which to look for the self-signed certificates to convert. (String, optional)

rootCertificateAlias

Specifies the root certificate to use from the default root store used to sign the chained certificate. The default value is `root`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask convertSelfSignedCertificatesToChained {-certificateReplacementOption ALL_CERTIFICATES -keyStoreName testKS}
```

- Using Jython string:

```
AdminTask.convertSelfSignedCertificatesToChained('[-certificateReplacementOption ALL_CERTIFICATES -keyStoreName testKS]')
```

- Using Jython list:

```
AdminTask.convertSelfSignedCertificatesToChained(['-certificateReplacementOption', 'ALL_CERTIFICATES', '-keyStoreName', 'testKS'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask exchangeSigners {-interactive}
```

- Using Jython:

```
AdminTask.exchangeSigners('-interactive')
```

enableWritableKeyrings command

The `enableWritableKeyrings` command modifies the keystore and enables writable SAF support. The system uses this command during migration. The command creates additional writable keystore objects for the control region and servant region key rings for SSL keystores.

Required parameters

-keyStoreName

Specifies the name that uniquely identifies the keystore that you want to delete. (String, required)

Optional parameters

-controlRegionUser

Specifies the control region user to use to enable writable key rings. (String, optional)

-servantRegionUser

Specifies the servant region user to enable writable key rings. (String, optional)

-scopeName

Specifies the name that uniquely identifies the management scope, for example: `(cell):localhostNode01Cell`. (String, optional)

Examples

Batch mode example usage:

- Using Jython string:

```
AdminTask.enableWritableKeyrings(['-keyStoreName testKS -controlRegionUser CRUser1 -servantRegionUser SRUser1'])
```

- Using Jython list:

```
AdminTask.enableWritableKeyrings(['-keyStoreName', 'testKS', '-controlRegionUser', 'CRUser1', '-servantRegionUser', 'SRUser1'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.enableWritableKeyrings('-interactive')
```

convertSSLConfig command

The convertSSLConfig command migrates existing SSL configurations to the new configuration object format for SSL configurations.

Required parameters

-sslConversionOption

Specifies how the system converts the SSL configuration. Specify the CONVERT_SSLCONFIGS value to convert the SSL configuration objects from the previous SSL configuration object to the new SSL configuration object. Specify the CONVERT_TO_DEFAULT value to convert the SSL configuration to a centralized SSL configuration, which also removes the SSL configuration direct referencing from the servers.

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jython string:

```
AdminTask.convertSSLConfig(['-keyStoreName testKS -controlRegionUser CRUser1 -servantRegionUser SRUser1'])
```

- Using Jython list:

```
AdminTask.convertSSLConfig(['-keyStoreName', 'testKS', '-controlRegionUser', 'CRUser1', '-servantRegionUser', 'SRUser1'])
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.convertSSLConfig('-interactive')
```

Related concepts

“Key management for cryptographic uses” on page 716

WebSphere Application Server provides a framework for managing keys (secret keys or key pairs) that applications use to perform cryptographic operations on data. The key management framework provides an application programming interface (API) for retrieving these keys. Keys are managed in keystores so the keystore type can be supported by WebSphere Application Server, provided that the keystores can store the referenced key type. You can configure keys and scope keystores so that they are visible only to particular processes, nodes, clusters, and so on.

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

“Automating SSL configurations using scripting” on page 960

SSL configuration is needed for WebSphere to perform SSL connections with other servers. A SSL configuration can be configured through the Admin Console. But if an automated way to create a SSL configuration is desired then AdminTask should be used.

“Creating an SSL configuration at the node scope using scripting” on page 957

An Secure Socket Layer (SSL) configuration references many other configuration objects. To help you make valid selections for the new SSL configuration before you create it, view information about existing configuration objects. Information about existing objects is also useful when you create a node scoped SSL configuration using the **createSSLConfig** command of the AdminTask object.

Related reference

“KeyStoreCommands command group for the AdminTask object” on page 981

You can use the Jython or Jacl scripting languages to configure keystores with the wsadmin tool. A keystore is created by the application server during install and can contain cryptographic keys or certificates. The commands and parameters in the KeyStoreCommands group can be used to create, delete, and manage keystores.

IdMgrConfig command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure the virtual member manager with the wsadmin tool. The commands and parameters in the IdMgrConfig group can be used to create and manage your entity type configuration.

The IdMgrConfig command group for the AdminTask object includes the following commands:

- “deleteIdMgrSupportedEntityType”
- “getIdMgrSupportedEntityType” on page 1159
- “listIdMgrSupportedEntityTypes” on page 1160
- “resetIdMgrConfig” on page 1160
- “showIdMgrConfig” on page 1161
- “updateIdMgrLDAPBindInfo” on page 1161
- “updateIdMgrSupportedEntityType” on page 1162

deleteIdMgrSupportedEntityType

The **deleteIdMgr Supported EntityType** command deletes the supported entity type configuration that you specify.

Parameters and return values

-name

The name of the supported entity type. The value of this parameter must be one of the supported entity types. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteIdMgrSupported EntityType {-name entity1}`
- Using Jython string:
`AdminTask.deleteIdMgrSupported EntityType ('[-name entity1']')`
- Using Jython list:
`AdminTask.deleteIdMgr SupportedEntityType (['-name', 'entity1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteIdMgr SupportedEntityType {-interactive}`
- Using Jython string:
`AdminTask.deleteIdMgr SupportedEntityType ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteIdMgr SupportedEntityType (['-interactive'])`

getIdMgrSupportedEntityType

The **getIdMgr Supported EntityType** command returns the configuration of the supported entity type that you specify.

Parameters and return values

-name

The name of the supported entity type. The value of this parameter must be one of the supported entity types. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask getIdMgrSupported EntityType {-name entity1}`
- Using Jython string:
`AdminTask.getIdMgrSupported EntityType ('[-name entity1']')`
- Using Jython list:
`AdminTask.getIdMgrSupported EntityType (['-name', 'entity1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask getIdMgrSupported EntityType {-interactive}`
- Using Jython string:
`AdminTask.getIdMgrSupported EntityType ('[-interactive]')`
- Using Jython list:
`AdminTask.getIdMgrSupported EntityType (['-interactive'])`

listIdMgrSupportedEntityTypes

The **listIdMgr Supported EntityTypes** command lists all of the supported entity types that are configured.

Parameters and return values

- Parameters: None
- Returns: A list that contains the names of the supported entity types.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listIdMgr SupportedEntityTypes`
- Using Jython string:
`AdminTask.listIdMgr SupportedEntityTypes()`
- Using Jython list:
`AdminTask.listIdMgr SupportedEntityTypes()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listIdMgrSupported EntityTypes {-interactive}`
- Using Jython string:
`AdminTask.listIdMgrSupported EntityTypes ('[-interactive]')`
- Using Jython list:
`AdminTask.listIdMgrSupported EntityTypes (['-interactive'])`

resetIdMgrConfig

The **resetId MgrConfig** command resets the current configuration to the last configuration that was saved.

Parameters and return values

- Parameters: None
- Returns: None

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask resetIdMgrConfig`
- Using Jython string:
`AdminTask.resetIdMgrConfig()`
- Using Jython list:
`AdminTask.resetIdMgrConfig()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask resetIdMgrConfig {-interactive}`
- Using Jython string:
`AdminTask.resetIdMgrConfig ('[-interactive]')`
- Using Jython list:

```
AdminTask.resetIdMgrConfig (['-interactive'])
```

showIdMgrConfig

The **showIdMgrConfig** command returns the current configuration XML in string format.

Parameters and return values

- Parameters: None
- Returns: None

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask showIdMgrConfig
```
- Using Jython string:

```
AdminTask.showIdMgrConfig()
```
- Using Jython list:

```
AdminTask.showIdMgrConfig()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask showIdMgrConfig {-interactive}
```
- Using Jython string:

```
AdminTask.showIdMgrConfig (['-interactive'])
```
- Using Jython list:

```
AdminTask.showIdMgrConfig (['-interactive'])
```

updateIdMgrLDAPBindInfo

The **updateIdMgrLDAPBindInfo** command updates the LDAP password under the federated repository. The change is also reflected in the wimconfig.xml file.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jython:

```
AdminTask.updateIdMgrLDAPBindInfo(['-id id1 -bindDN cn=root -bindPassword myPassword22'])
```
- Using Jython list:

```
AdminTask.updateIdMgrLDAPBindInfo(['-id id1 -bindDN cn=root -bindPassword myPassword22'])
```
- Using Jacl:

```
$AdminTask updateIdMgrLDAPBindInfo {-id id1 -bindDN cn=root -bindPassword myPassword22}
```

Interactive mode example usage:

- Using Jython:

```
AdminTask.updateIdMgrLDAPBindInfo(['-interactive'])
```
- Using Jacl:


```
$AdminTask updateIdMgrLDAPBindInfo {-interactive}
```

updateIdMgrSupportedEntityType

The **updateIdMgr Supported EntityType** command updates the configuration that you specify for a supported entity type.

Parameters and return values

-name

The name of the supported entity type. The value of this parameter must be one of the supported entity types. (String, required)

-defaultParent

The default parent node for the supported entity type. (String, optional)

-rdnProperties

The RDN attribute name for the supported entity type in the entity domain name. To reset all the values of the `rdnProperties` parameter, specify a blank string (`""`). (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrSupported EntityType {-name entity1}
```
- Using Jython string:

```
AdminTask.updateIdMgrSupported EntityType ('[-name entity1']')
```
- Using Jython list:

```
AdminTask.updateIdMgrSupported EntityType (['-name', 'entity1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrSupported EntityType {-interactive}
```
- Using Jython string:

```
AdminTask.updateIdMgrSupported EntityType ('[-interactive]')
```
- Using Jython list:

```
AdminTask.updateIdMgrSupported EntityType (['-interactive'])
```

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

“IdMgrRepositoryConfig command group for the AdminTask object”

You can use the Jython or Jacl scripting languages to configure security. The commands and parameters in the IdMgrRepositoryConfig group can be used to create and manage the virtual member manager and LDAP directory properties.

“IdMgrRealmConfig command group for the AdminTask object” on page 1210

You can use the Jython or Jacl scripting languages to configure the member manager realm and realms. The commands and parameters in the IdMgrRealmConfig group can be used to create and manage your realm configuration.

IdMgrRepositoryConfig command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security. The commands and parameters in the IdMgrRepositoryConfig group can be used to create and manage the virtual member manager and LDAP directory properties.

The IdMgrRepositoryConfig command group for the AdminTask object includes the following commands:

- “addIdMgrLDAPAttr” on page 1164
- “addIdMgrLDAPBackupServer” on page 1165
- “addIdMgrLDAPEntityType” on page 1166
- “addIdMgrLDAPEntityTypeRDNAttr” on page 1167
- “addIdMgrLDAPGroupDynamicMemberAttr” on page 1167
- “addIdMgrLDAPGroupMemberAttr” on page 1168
- “addIdMgrLDAPServer” on page 1169
- “addIdMgrRepositoryBaseEntry” on page 1170
- “createIdMgrCustomRepository” on page 1171
- “createIdMgrDBRepository” on page 1171
- “createIdMgrFileRepository” on page 1173
- “createIdMgrLDAPRepository” on page 1174
- “deleteIdMgrLDAPAttr” on page 1175
- “deleteIdMgrLDAPEntityType” on page 1176
- “deleteIdMgrLDAPEntityTypeRDNAttr” on page 1176
- “deleteIdMgrLDAPGroupConfig” on page 1177
- “deleteIdMgrLDAPGroupMemberAttr” on page 1177
- “deleteIdMgrLDAPGroupDynamicMemberAttr” on page 1178
- “deleteIdMgrLDAPServer” on page 1178
- “deleteIdMgrRepository” on page 1179
- “deleteIdMgrRepositoryBaseEntry” on page 1179
- “getIdMgrLDAPAttrCache” on page 1180
- “getIdMgrLDAPContextPool” on page 1180
- “getIdMgrLDAPEntityType” on page 1181
- “getIdMgrLDAPEntityTypeRDNAttr” on page 1181

- “getldMgrLDAPGroupConfig” on page 1182
- “getldMgrLDAPGroupDynamicMemberAttrs” on page 1182
- “getldMgrLDAPGroupMemberAttrs” on page 1183
- “getldMgrLDAPSearchResultCache” on page 1183
- “getldMgrLDAPServer” on page 1184
- “getldMgrRepository” on page 1184
- “listldMgrLDAPAttrs” on page 1185
- “listldMgrCustomProperties” on page 1185
- “listldMgrLDAPBackupServers” on page 1186
- “listldMgrLDAPEntityTypes” on page 1186
- “listldMgrLDAPServers” on page 1187
- “listldMgrRepositories” on page 1187
- “listldMgrRepositoryBaseEntries” on page 1188
- “listldMgrSupportedDBTypes” on page 1189
- “listldMgrSupportedMessageDigestAlgorithms” on page 1189
- “listldMgrSupportedLDAPServerTypes” on page 1190
- “removeIdMgrLDAPBackupServer” on page 1190
- “setldMgrCustomProperty” on page 1191
- “setldMgrLDAPAttrCache” on page 1191
- “setldMgrLDAPContextPool” on page 1193
- “setldMgrLDAPGroupConfig” on page 1194
- “setldMgrLDAPSearchResultCache” on page 1195
- “setldMgrEntryMappingRepository” on page 1196
- “setldMgrPropertyExtensionRepository” on page 1197
- “updateIdMgrDBRepository” on page 1198
- “updateIdMgrFileRepository” on page 1198
- “updateIdMgrLDAPAttrCache” on page 1199
- “updateIdMgrLDAPContextPool” on page 1201
- “updateIdMgrLDAPEntityType” on page 1202
- “updateIdMgrLDAPGroupDynamicMemberAttr” on page 1203
- “updateIdMgrLDAPGroupMemberAttr” on page 1203
- “updateIdMgrLDAPRepository” on page 1204
- “updateIdMgrLDAPSearchResultCache” on page 1206
- “updateIdMgrLDAPServer” on page 1207
- “updateIdMgrRepository” on page 1208
- “updateIdMgrRepositoryBaseEntry” on page 1209

addldMgrLDAPAttr

The addldMgrLDAPAttr command adds an LDAP attribute configuration to the LDAP repository configuration.

Required parameters

-id Specifies the unique ID of the repository. (String, required)

-name

Specifies the name of the LDAP attribute used in the repository LDAP adapter. (String, required)

Optional parameters

-entityTypes

Specifies the entity type which applies the attribute mapping. (String, optional)

-syntax

Specifies the syntax of the LDAP attribute. The default value is `string`. For example, the syntax of the `unicodePwd` LDAP attribute is `octetString`. (String, optional)

-defaultValue

Specifies the default value of the LDAP attribute. If you do not specify this LDAP attribute when you create an entity which this LDAP attribute applies to, the system adds the attribute using this default value. (String, optional)

-defaultAttr

The default attribute of the LDAP attribute. If you do not specify this LDAP attribute when you create an entity which this LDAP attribute applies to, the system uses this value of the default attribute. (String, optional)

-propertyName

Specifies the name of the corresponding federated repository property. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPAttr {-id id1 -name unicodePwd -syntax octetString}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPAttr ('[-id id1 -name unicodePwd -syntax octetString']')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPAttr (['-id', 'id1', '-name', 'unicodePwd', '-syntax', 'octetString'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPAttr {-interactive}
```

- Using Jython:

```
AdminTask.addIdMgrLDAPAttr('-interactive')
```

addIdMgrLDAPBackupServer

The **addIdMgrLDAPBackupServer** command sets a backup LDAP server in your configuration.

Required parameters

-id Specifies the unique ID of the repository. (String, required)

-primary_host

Specifies the primary host of the LDAP server. (String, required)

-host

Specifies the host name for the LDAP server. (String, required)

Optional parameters

-port

Specifies the port number for the LDAP server. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPBackupServer {-id id1 -primary_host hostName -host hostName2 -port 4020}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPBackupServer ('[-id id1 -primary_host hostName -host hostName2 -port 4020']')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPBackupServer (['-id', 'id1', '-primary_host', 'hostName', '-host', 'hostName2', '-port', '4020'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPBackupServer {-interactive}
```

- Using Jython:

```
AdminTask.addIdMgrLDAPBackupServer('-interactive')
```

addIdMgrLDAPEntityType

The **addIdMgrLDAPEntityType** command adds an LDAP entity type definition.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the entity type. (String, required)

-searchFilter

The search filter that you want to use to search the entity type. (String, optional)

-objectClasses

One or more object classes for the entity type. (String, required)

-objectClassesForCreate

The object class to use when an entity type is created. If the value of this parameter is the same as the objectClass parameter, you do not need to specify this parameter. (String, optional)

-searchBases

The search base or bases to use while searching the entity type. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPEntityType {-id id1 -name name1 -objectClasses objectclass}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPEntityType ('[-id id1 -name name1 -objectClasses objectclass']')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPEntityType (['-id', 'id1', '-name', 'name1', '-objectClasses', 'objectclass'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPEntityType {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPEntityType ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPEntityType (['-interactive'])
```

addIdMgrLDAPEntityTypeRDNAttr

The **addId MgrLDAP EntityType RDNAttr** command adds RDN attribute configuration to an LDAP entity type definition.

Parameters and return values

-id The ID of the repository. (String, required)

-entityTypeName

The name of the entity type. (String, required)

-name

The attribute name that is used to build the relative distinguished name (RDN) for the entity type. (String, required)

-objectClass

The object class to use for the entity type for the relative distinguished name (RDN) attribute name that you specify. Use this parameter to map one entity type to multiple structural object classes. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPEntityTypeRDNAttr {-id id1 -entityTypeName entitytype -name name1}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPEntityTypeRDNAttr ('[-id id1 -entityTypeName entitytype -name name1]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPEntityTypeRDNAttr (['-id', 'id1', '-entityTypeName', 'entity type', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPEntityTypeRDNAttr {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPEntityTypeRDNAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPEntityTypeRDNAttr (['-interactive'])
```

addIdMgrLDAPGroupDynamicMemberAttr

The **addIdMgr LDAPGroup Dynamic Member Attr** command adds a dynamic member attribute configuration to an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the LDAP attribute that is used as the group member attribute. For example, member or uniqueMember. (String, required)

-objectClass

The group object class that contains the member attribute. For example, groupOfNames or groupOfUniqueNames. If you do not define this parameter, the member attribute applies to all group object classes. (String, optional)

-scope

The scope of the member attribute. The valid values for this parameter include the following:

- **direct** - The member attribute only contains direct members, therefore, this value refers to the member directly contained by the group and not contained through the nested group. For example, if Group1 contains Group2 and Group2 contains User1, then Group2 is a direct member of Group1 but User1 is not a direct member of Group1. Both `member` and `uniqueMember` are direct member attributes.
- **nested** - The member attribute that contains the direct members and the nested members.

-dummyMember

Indicates that if you create a group without specifying a member, a dummy member will be filled in to avoid creating an exception about missing a mandatory attribute. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPGroup DynamicMemberAttr {-id id1 -name name1 -objectClass objectclass}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPGroup DynamicMemberAttr ('[-id id1 -name name1 -objectClass objectclass']')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPGroup DynamicMemberAttr (['-id', 'id1', '-name', 'name1', '-objectClass', 'objectclass'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPGroup DynamicMemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPGroup DynamicMemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPGroup DynamicMemberAttr (['-interactive'])
```

addIdMgrLDAPGroupMemberAttr

The **addIdMgr LDAPGroup MemberAttr** command adds a member attribute configuration to an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the LDAP attribute that is used as the group member attribute. For example, `member` or `uniqueMember`. (String, required)

-objectClass

The group object class that contains the member attribute. For example, `groupOfNames` or `groupOfUniqueNames`. If you do not define this parameter, the member attribute applies to all group object classes. (String, optional)

-scope

The scope of the member attribute. The valid values for this parameter include the following:

- **direct** - The member attribute only contains direct members, therefore, this value refers to the member directly contained by the group and not contained through the nested group. For example, if Group1 contains Group2 and Group2 contains User1, then Group2 is a direct member of Group1 but User1 is not a direct member of Group1. Both `member` and `uniqueMember` are direct member attributes.

- nested - The member attribute that contains the direct members and the nested members.

-dummyMember

Indicates that if you create a group without specifying a member, a dummy member will be filled in to avoid creating an exception about missing a mandatory attribute. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPGroup MemberAttr {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPGroup MemberAttr ('[-id id1 -name name1]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPGroup MemberAttr (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPGroup MemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPGroup MemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPGroup MemberAttr (['-interactive'])
```

addIdMgrLDAPServer

The **addId MgrLDAP Server** command adds an LDAP server to the LDAP repository ID that you specify.

Parameters and return values

-id The ID of the repository. (String, required)

-host

The host name for the primary LDAP server. (String, required)

-port

The port number for the LDAP server. (Integer, optional)

-bindDN

The binding distinguished name for the LDAP server. (String, optional)

-bindPassword

The binding password. (String, optional)

-authentication

Indicates the authentication method to use. The default value is `simple`. Valid values include: `none` or `strong`. (String, optional)

-referral

The LDAP referral. The default value is `ignore`. Valid values include: `follow`, `throw`, or `false`. (String, optional)

-derefAliases

Controls how aliases are dereferenced. The default value is `always`. Valid values include:

- `never` - never deference aliases
- `finding` - deferences aliases only during name resolution
- `searching` - deferences aliases only after name resolution

(String, optional)

-sslEnabled

Indicates to enable SSL or not. The default value is false. (Boolean, optional)

-connectionPool

The connection pool. The default value is false. (Boolean, optional)

-connectTimeout

The connection timeout in seconds. The default value is 0. (Integer, optional)

-ldapServerType

The type of LDAP server being used. The default value is IDS51. (String, optional)

-sslConfiguration

The SSL configuration. (String, optional)

-certificateMapMode

Specifies whether to map X.509 certificates into a LDAP directory by exact distinguished name or by certificate filter. The default value is EXACT_DN. To use the certificate filter for the mapping, specify CERTIFICATE_FILTER. (String, optional)

-certificateFilter

If certificateMapMode has the value CERTIFICATE_FILTER, then this property specifies the LDAP filter which maps attributes in the client certificate to entries in LDAP. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAPServer {-id id1 -host myhost.ibm.com}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAPServer ('[-id id1 -host myhost.ibm.com]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAPServer (['-id', 'id1', '-host', 'myhost.ibm.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrLDAP Server {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrLDAP Server ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrLDAP Server (['-interactive'])
```

addIdMgrRepositoryBaseEntry

The **addIdMgr Repository BaseEntry** command adds a base entry to the specified repository.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The distinguished name of a base entry. (String, required)

-nameInRepository

The distinguished name in the repository that uniquely identifies the base entry name. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrRepositoryBaseEntry {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.addIdMgrRepositoryBaseEntry ('[-id id1 -name name1]')
```

- Using Jython list:

```
AdminTask.addIdMgrRepositoryBaseEntry (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrRepositoryBaseEntry {-interactive}
```

- Using Jython string:

```
AdminTask.addIdMgrRepositoryBaseEntry ('[-interactive]')
```

- Using Jython list:

```
AdminTask.addIdMgrRepositoryBaseEntry (['-interactive'])
```

createIdMgrCustomRepository

The **createIdMgrCustomRepository** command creates a custom repository configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-adapterClassName

The implementation class name for the repository adapter. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrCustomRepository {-id id1 -adapterClassName adapterClassName}
```

- Using Jython string:

```
AdminTask.createIdMgrCustomRepository('-id id1 -adapterClassName adapterClassName')
```

- Using Jython list:

```
AdminTask.createIdMgrCustomRepository(['-id', 'id1', '-adapterClassName', 'adapterClassName'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrCustomRepository {-interactive}
```

- Using Jython:

```
AdminTask.createIdMgrCustomRepository('-interactive')
```

createIdMgrDBRepository

The **createIdMgrDBRepository** command creates a database repository configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-dataSourceName

The name of the data source. The default value is jdbc/wimDS. (String, required)

-databaseType

The type of the database. The default value is DB2. (String, required)

-dbURL

The URL of the database. (String, required)

-dbAdminId

The database administrator ID. (String, required if database type is not Apache Derby.)

-dbAdminPassword

The database administrator password. (String, required if database type is not Apache Derby.)

-adapterClassName

The default value is com.ibm.ws.wim.adapter.db.DBAdapter. (String, optional)

-JDBCDriverClass

The JDBC driver class name. (String, optional)

-supportSorting

Indicates if sorting is supported or not. The default value is false. (Boolean, optional)

-supportTransactions

Indicates if transactions are supported or not. The default value is false. (Boolean, optional)

-isExtIdUnique

Specifies if the external ID is unique. The default value is true. (Boolean, optional)

-supportExternalName

Indicates if external names are supported or not. The default value is false. (Boolean, optional)

-entityRetrievalLimit

Indicates the value of the retrieval limit on database entries. The default value is 200. (Integer, optional)

-saltLength

The salt length in bits. The default value is 12. (Integer, optional)

-encryptionKey

The default value is rZ15ws0e1y9yHk3zCs3sTMv/ho8fY17s. (String, optional)

Examples**Batch mode example usage:**

- Using Jacl:

```
$AdminTask createIdMgrDB Repository {-id id1 -dataSourceName datasource name -databaseType DB2}
```

- Using Jython string:

```
AdminTask.createIdMgrDB Repository ('[-id id1 -dataSourceName data sourcename -databaseType DB2']')
```

- Using Jython list:

```
AdminTask.createIdMgrDB Repository (['-id', 'id1', '-dataSourceName', 'data sourcename', '-databaseType', 'DB2'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrDB Repository {-interactive}
```

- Using Jython string:

```
AdminTask.createIdMgrDB Repository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.createIdMgrDB Repository (['-interactive'])
```

createIdMgrFileRepository

The **createId MgrFile Repository** command creates a file repository configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-messageDigest Algorithm

The message digest algorithm that will be used for hashing the password. The default value is SHA-1. Valid values include the following: SHA-1, SHA-384, or SHA-512.(String, required)

-adapterClassName

The default value is `com.ibm.ws.wim.adapter.file.was.FileAdapter`. (String, optional)

-supportPaging

Indicates if paging is supported or not. The default value is `false`. (Boolean, optional)

-supportSorting

Indicates if sorting is supported or not. The default value is `false`. (Boolean, optional)

-supportTransactions

Indicates if transaction is supported or not. The default value is `false`. (Boolean, optional)

-isExtIdUnique

Specifies if the external ID is unique or not. The default value is `true`. (Boolean, optional)

-supportExternalName

Indicates if external names are supported or not. The default value is `false`. (Boolean, optional)

-baseDirectory

The base directory where the file will be created in order to store the data. The default is to be dynamically built during run time using `user.install.root` and cell name. (String, optional)

-fileName

The file name of the repository. The default value is `fileRegistry.xml`. (String, optional)

-saltLength

The salt length of the randomly generated salt for password hashing. The default value is 12. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrFile Repository {-id id1 -messageDigestAlgorithm SHA-1}
```

- Using Jython string:

```
AdminTask.createIdMgrFile Repository (['-id id1 -messageDigestAlgorithm SHA-1'])
```

- Using Jython list:

```
AdminTask.createIdMgrFile Repository (['-id', 'id1', '-messageDigestAlgorithm', 'SHA-1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrFile Repository {-interactive}
```

- Using Jython string:

```
AdminTask.createIdMgrFile Repository (['-interactive'])
```

- Using Jython list:

```
AdminTask.createIdMgrFile Repository (['-interactive'])
```

createIdMgrLDAPRepository

The **create IdMgrLDAP Repository** command creates an LDAP repository configuration.

Parameters and return values

- id** The unique identifier for the repository. (String, required)
- ldapServerType**
The type of LDAP server that is being used. The default value is `IDS51`. (String, required)
- adapterClassName**
The default value is `com.ibm.ws.wim.adapter.db.DBAdapter`. (String, optional)
- supportSorting**
Indicates if sorting is supported or not. The default value is `false`. (Boolean, optional)
- supportPaging**
Indicates if paging is supported or not. The default value is `false`. (Boolean, optional)
- supportTransactions**
Indicates if transactions are supported or not. The default value is `false`. (Boolean, optional)
- isExtIdUnique**
Specifies if the external ID is unique. The default value is `true`. (Boolean, optional)
- supportExternalName**
Indicates if external names are supported or not. The default value is `false`. (Boolean, optional)
- authentication**
Indicates the authentication method to use. The default value is `simple`. Valid values include: `none` or `strong`. (String, optional)
- referral**
The LDAP referral. The default value is `ignore`. Valid values include: `follow`, `throw`, or `false`. (String, optional)
- sslEnabled**
Indicates to enable SSL or not. The default value is `false`. (Boolean, optional)
- sslConfiguration**
The SSL configuration. (String, optional)
- connectionPool**
The connection pool. The default value is `false`. (Boolean, optional)
- translateRDN**
Indicates to translate RDN or not. The default value is `false`. (Boolean, optional)
- searchTimeLimit**
The value of search time limit. (Integer, optional)
- searchCountLimit**
The value of search count limit. (Integer, optional)
- searchPageSize**
The value of search page size. (Integer, optional)
- returnToPrimaryServer**
(Integer, optional)
- primaryServerQueryTimeInterval**
(Integer, optional)

-default

If you set this parameter to true, the default values will be set for the remaining configuration properties of the LDAP repository. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrLDAP Repository {-id id1 -ldapServerType IDS51}
```

- Using Jython string:

```
AdminTask.createIdMgrLDAP Repository ('[-id id1 -ldapServerType IDS51]')
```

- Using Jython list:

```
AdminTask.createIdMgrLDAP Repository (['-id', 'id1', '-ldapServerType', 'IDS51'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrLDAP Repository {-interactive}
```

- Using Jython string:

```
AdminTask.createIdMgrLDAP Repository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.createIdMgrLDAP Repository (['-interactive'])
```

deleteIdMgrLDAPAttr

The deleteIdMgrLDAPAttr command deletes LDAP attribute configuration data for a specific entity type from the LDAP repository of interest.

Required parameters

-id Specifies the unique ID of the repository. (String, required)

-name

Specifies the name of the LDAP attribute used in the repository LDAP adapter. (String, required)

Optional parameters

-entityTypes

Specifies the entity type which applies the attribute mapping. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAPAttr {-id id1 -name unicodePwd -syntax octetString}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAPAttr ('[-id id1 -name unicodePwd -syntax octetString]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAPAttr (['-id', 'id1', '-name', 'unicodePwd', '-syntax', 'octetString'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAPAttr {-interactive}
```

- Using Jython:

```
AdminTask.deleteIdMgrLDAPAttr('-interactive')
```

deleteIdMgrLDAPEntityType

The **deleteId MgrLDAP EntityType** command deletes the LDAP entity type configuration data for a specified entity type for a specific LDAP repository.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the entity type. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP EntityType {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP EntityType ('[-id id1 -name name1]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP EntityType (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP EntityType {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP EntityType ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP EntityType (['-interactive'])
```

deleteIdMgrLDAPEntityTypeRDNAttr

The **deleteId MgrLDAP EntityType RDNAttr** command deletes the relative distinguished name (RDN) attribute configuration from an LDAP entity type configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-entityTypeName

The name of the entity type. (String, required)

-name

The attribute name that is used to build the relative distinguished name (RDN) for the entity type. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAPEntity TypeRDNAttr {-id id1 -name name1 -entityTypeName entityType}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP Entity TypeRDNAttr ('[-id id1 -name name1 -entityType Name entityType]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAPEntity TypeRDNAttr (['-id', 'id1', '-name', 'name1', '-entity TypeName', 'entityType'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAPEntity TypeRDNAttr {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAPEntity TypeRDNAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAPEntity TypeRDNAttr (['-interactive'])
```

deleteIdMgrLDAPGroupConfig

The **deleteIdMgrLDAP GroupConfig** command deletes the LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP GroupConfig {-id id1}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP GroupConfig ('[-id id1]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP GroupConfig (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP GroupConfig {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP GroupConfig ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP GroupConfig (['-interactive'])
```

deleteIdMgrLDAPGroupMemberAttr

The **deleteIdMgr LDAPGroup MemberAttr** command deletes a member attribute configuration from an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP GroupMemberAttr {-id id1}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP GroupMemberAttr ('[-id id1]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP GroupMemberAttr (['-id', 'id1'])
```


Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP GroupMemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP GroupMemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP GroupMemberAttr (['-interactive'])
```

deleteIdMgrLDAPGroupDynamicMemberAttr

The **deleteIdMgr LDAPGroup Dynamic MemberAttr** command deletes a dynamic member attribute configuration from an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the LDAP attribute that is used as the group member attribute. For example, memberURL. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP GroupDynamicMemberAttr {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP GroupDynamicMemberAttr ('[-id id1 -name name1]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP GroupDynamicMemberAttr (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAPGroup DynamicMemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAPGroup DynamicMemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAPGroup DynamicMemberAttr (['-interactive'])
```

deleteIdMgrLDAPServer

The **deleteId MgrLDAP Server** command deletes the configuration for the LDAP server that you specify from the LDAP repository ID that you specify.

Parameters and return values

-id The ID of the repository. (String, required)

-host

The host name for the primary LDAP server. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP Server {-id id1 -host myhost.ibm.com}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP Server ('[-id id1 -host myhost.ibm.com]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP Server (['-id', 'id1', '-host', 'myhost.ibm.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrLDAP Server {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrLDAP Server ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrLDAP Server (['-interactive'])
```

deleteIdMgrRepository

The **deleteIdMgr Repository** command deletes a repository that you specify.

Parameters and return values

-id The ID of the repository. Valid values include existing repository IDs. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgr Repository {-id id1}
```

- Using Jython string:

```
AdminTask.deleteIdMgr Repository ('[-id id1]')
```

- Using Jython list:

```
AdminTask.deleteIdMgr Repository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrRepository {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrRepository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrRepository (['-interactive'])
```

deleteIdMgrRepositoryBaseEntry

The **deleteIdMgr Repository BaseEntry** command deletes a base entry from the specified repository.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The distinguished name of a base entry. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrRepository BaseEntry {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.deleteIdMgrRepository BaseEntry ('[-id id1 -name name1']')
```

- Using Jython list:

```
AdminTask.deleteIdMgrRepository BaseEntry (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrRepository BaseEntry {-interactive}
```

- Using Jython string:

```
AdminTask.deleteIdMgrRepository BaseEntry ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteIdMgrRepository BaseEntry (['-interactive'])
```

getIdMgrLDAPAttrCache

The **getIdMgr LDAPAttr Cache** command returns the LDAP attribute cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPAttr Cache {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPAttr Cache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPAttr Cache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAP AttrCache {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPAttr Cache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPAttr Cache (['-interactive'])
```

getIdMgrLDAPContextPool

The **getIdMgr LDAP Context Pool** command returns the LDAP context pool configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPContext Pool {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPContext Pool ('[-id id1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPContext Pool (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPCon textPool {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPCon textPool ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPCon textPool (['-interactive'])
```

getIdMgrLDAPEntityType

The **getIdMgr LDAP EntityType** command returns the LDAP entity type configuration data.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the entity type. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPEntity Type {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPEntity Type ('[-id id1 -name name1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPEntity Type (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPEn tityType {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPEn tityType ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPEn tityType (['-interactive'])
```

getIdMgrLDAPEntityTypeRDNAttr

The **getIdMgr LDAPEntity TypeRDNAttr** command returns the relative distinguished name (RDN) attribute configuration for an LDAP entity type definition.

Parameters and return values

-id The ID of the repository. (String, required)

-entityTypeName

The name of the entity name. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPEntity TypeRDNAttr {-id id1 -entityTypeName name1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPEntity TypeRDNAttr ('[-id id1 -entityTypeName name1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPEntity TypeRDNAttr (['-id', 'id1', '-entityTypeName', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPEntity TypeRDNAttr {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPEntity TypeRDNAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPEntity TypeRDNAttr (['-interactive'])
```

getIdMgrLDAPGroupConfig

The **getIdMgr LDAPG roupConfig** command returns the LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPGroup Config {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPGroup Config ('[-id id1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPGroup Config (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPGroup Config {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPGroup Config ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPGroup Config (['-interactive'])
```

getIdMgrLDAPGroupDynamicMemberAttrs

The **getIdMgr LDAPGroup Dynamic Member Attrs** command returns the dynamic member attribute configuration from the LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPGroupDynamic MemberAttrs {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPGroupDynamic MemberAttrs ('[-id id1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPGroupDynamic MemberAttrs (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPGroup DynamicMemberAttrs {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPGroup DynamicMemberAttrs ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPGroup DynamicMemberAttrs (['-interactive'])
```

getIdMgrLDAPGroupMemberAttrs

The **getIdMgr LDAPGroup MemberAttrs** command returns the member attribute configuration for the LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPGroup MemberAttrs {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPGroup MemberAttrs ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPGroup MemberAttrs (['-interactive'])
```

getIdMgrLDAPSearchResultCache

The **getIdMgr LDAPSearch ResultCache** command returns the LDAP search result cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAP SearchResultCache {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAP SearchResultCache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPSearchResultCache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPSearchResultCache {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPSearchResultCache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPSearchResultCache (['-interactive'])
```

getIdMgrLDAPServer

The **getIdMgr LDAPServer** command returns the configuration for the LDAP server that you specify for the LDAP repository ID that you specify.

Parameters and return values

-id The ID of the repository. (String, required)

-host

The host name for the primary LDAP server. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPServer {-id id1 -host myhost.ibm.com}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPServer ('[-id id1 -host myhost.ibm.com]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPServer (['-id', 'id1', '-host', 'myhost.ibm.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrLDAPServer {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrLDAPServer ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrLDAPServer (['-interactive'])
```

getIdMgrRepository

The **getIdMgr Repository** command returns the configuration of the specified repository.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrRepository {-id id1}
```

- Using Jython string:

```
AdminTask.getIdMgrRepository ('[-id id1]')
```

- Using Jython list:

```
AdminTask.getIdMgrRepository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrRepository {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrRepository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getIdMgrRepository (['-interactive'])
```

listIdMgrLDAPAttrs

The `listIdMgrLDAPAttrs` command lists the name of each configured attributes for the LDAP repository of interest.

Required parameters

-id Specifies the unique ID of the repository. (String, required)

Return value

The command returns a list of HashMaps that contains parameters of the `addIdMgrLDAPAttr` command as keys.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAPAttrs {-id id1}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAPAttrs ('[-id id1]')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAPAttrs (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAPAttrs {-interactive}
```

- Using Jython:

```
AdminTask.listIdMgrLDAPAttrs('-interactive')
```

listIdMgrCustomProperties

The `listIdMgr Custom Properties` command returns a list of custom properties for the repository that you specify.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrCustom Properties {-id id1}
```

- Using Jython string:

```
AdminTask.listIdMgrCustom Properties ('[-id id1']')
```

- Using Jython list:

```
AdminTask.listIdMgrCustom Properties (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrCustom Properties {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrCustom Properties ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listIdMgrCustom Properties (['-interactive'])
```

listIdMgrLDAPBackupServers

The **listIdMgr LDAPBackupServers** command returns a list of the backup LDAP server or servers.

Parameters and return values

-id The ID of the repository. (String, required)

-primary_host

The host name for the primary LDAP server. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP BackupServer {-id id1 -primary_host hostname}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP BackupServer ('[-id id1 -primary_host hostname']')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAP BackupServer (['-id', 'id1', '-primary_host', 'hostname'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP BackupServer {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP BackupServer ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAP BackupServer (['-interactive'])
```

listIdMgrLDAPEntityTypeTypes

The **listIdMgr LDAPEntityTypeTypes** command lists the name of all of the configured LDAP entity type definitions.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP EntityType {-id id1}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP EntityType ('[-id id1']')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAPEntity Type (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP EntityType {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP EntityType ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAP EntityType (['-interactive'])
```

listIdMgrLDAPServers

The **listIdMgr LDAP Servers** command lists all of the configured primary LDAP servers.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP Servers {-id id1}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP Servers ('[-id id1']')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAP Servers (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrLDAP Servers {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrLDAP Servers ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listIdMgrLDAP Servers (['-interactive'])
```

listIdMgrRepositories

The **listIdMgr Repositories** command lists names and types of all configured repositories.

Parameters and return values

- Parameters: None
- Returns: A hash map with key as the name of the repository and value as another hash map that includes the following keys:

- repositoryType - The type of repository. For example, File, LDAP, DB, and so on.
- specificRepositoryType - The specific type of repository. For example, LDAP, IDS51, NDS, and so on.
- host - The host name where the repository resides. For File, it is LocalHost and for DB it is dataSourceName.

This command will not return the Property Extension and Entry Mapping repository data.

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask listIdMgrRepositories
```
- Using Jython string:


```
AdminTask.listIdMgrRepositories()
```
- Using Jython list:


```
AdminTask.listIdMgrRepositories()
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask listIdMgrRepositories {-interactive}
```
- Using Jython string:


```
AdminTask.listIdMgrRepositories (['-interactive'])
```
- Using Jython list:


```
AdminTask.listIdMgrRepositories (['-interactive'])
```

listIdMgrRepositoryBaseEntries

The **listIdMgr Repository BaseEntries** command lists the base entries for a specified repository.

Parameters and return values

-id The ID of the repository. (String, required)

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask listIdMgrRepository BaseEntries {-id id1}
```
- Using Jython string:


```
AdminTask.listIdMgrRepository BaseEntries (['-id id1'])
```
- Using Jython list:


```
AdminTask.listIdMgrRepository BaseEntries (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask listIdMgrRepository BaseEntries {-interactive}
```
- Using Jython string:


```
AdminTask.listIdMgrRepository BaseEntries (['-interactive'])
```
- Using Jython list:


```
AdminTask.listIdMgrRepository BaseEntries (['-interactive'])
```

listIdMgrSupportedDBTypes

The **listIdMgr Supported DBTypes** command returns a list of supported database types.

Parameters and return values

- Parameters: None
- Returns: A list of supported database types.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupportedDBTypes
```

- Using Jython string:

```
AdminTask.listIdMgrSupportedDBTypes()
```

- Using Jython list:

```
AdminTask.listIdMgrSupportedDBTypes()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupported DBTypes {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrSupported DBTypes (['-interactive'])
```

- Using Jython list:

```
AdminTask.listIdMgrSupported DBTypes (['-interactive'])
```

listIdMgrSupportedMessageDigestAlgorithms

The **listIdMgr Supported Message Digest Algorithms** command returns a list of supported message digest algorithms.

Parameters and return values

- Parameters: None
- Returns: A list of supported message digest algorithms.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupported MessageDigestAlgorithms
```

- Using Jython string:

```
AdminTask.listIdMgrSupported MessageDigestAlgorithms()
```

- Using Jython list:

```
AdminTask.listIdMgrSupported MessageDigestAlgorithms()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupportedMes sageDigestAlgorithms {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrSupportedMessage DigestAlgorithms (['-interactive'])
```

- Using Jython list:

```
AdminTask.listIdMgrSupportedMessage DigestAlgorithms (['-interactive'])
```

listIdMgrSupportedLDAPServerTypes

The **listIdMgr Supported LDAP ServerTypes** command returns a list of supported LDAP server types.

Parameters and return values

- Parameters: None
- Returns: A list of supported LDAP server types.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupported LDAPServerTypes
```
- Using Jython string:

```
AdminTask.listIdMgrSupported LDAPServerTypes()
```
- Using Jython list:

```
AdminTask.listIdMgrSupported LDAPServerTypes()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrSupported LDAPServerTypes {-interactive}
```
- Using Jython string:

```
AdminTask.listIdMgrSupported LDAPServerTypes (['-interactive'])
```
- Using Jython list:

```
AdminTask.listIdMgrSupported LDAPServerTypes (['-interactive'])
```

removeIdMgrLDAPBackupServer

The **removeIdMgr LDAPBack upServer** command removes the backup LDAP server or servers.

Parameters and return values

- id** The ID of the repository. (String, required)
- primary_host**
The host name for the primary LDAP server. (String, required)
- host**
The name of the backup host name. Use an asterisk (*) if you want to remove all backup servers. (String, required)
- port**
The port number of the LDAP server. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask removeIdMgrLDAP BackupServer {-id id1 -primary_host myprimaryhost.ibm.com -host myhost.ibm.com}
```
- Using Jython string:

```
AdminTask.removeIdMgrLDAPBackup Server (['-id id1 -primary_host myprimaryhost.ibm.com -host myhost.ibm.com'])
```
- Using Jython list:

```
AdminTask.removeIdMgrLDAPBack upServer (['-id', 'id1', '-primary_host', 'myprimary host.ibm.com', '-host', 'myhost.ibm.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask removeIdMgrLDAP BackupServer {-interactive}
```

- Using Jython string:

```
AdminTask.removeIdMgrLDAP BackupServer ('[-interactive]')
```

- Using Jython list:

```
AdminTask.removeIdMgrLDAP BackupServer (['-interactive'])
```

setIdMgrCustomProperty

The **setIdMgr Custom Property** command : sets, adds or deletes a custom property to a repository configuration. If a value is not specified, or if there is an empty string, the property is deleted from the repository configuration. If a name does not exist it is added if a value is specified. If the name is "*" then all of the custom properties are deleted.

Parameters and return values

-id The unique identifier of the repository. Valid values include the existing repository IDs. (String, required)

-name

The name of the additional property for the repository that are not defined OOTB.(String, required)

-value

The value of a property for the repository. If this parameter is an empty string, the property is deleted from the repository configuration. If this parameter is not an empty string, and a name does not exist, it is added. If a name is an empty string, all of the custom properties are deleted. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrCustomProperty {-id id1 -name name1 -value value}
```

- Using Jython string:

```
AdminTask.setIdMgrCustomProperty ('[-id id1 -name name1 -value value]')
```

- Using Jython list:

```
AdminTask.setIdMgrCustomProperty (['-id', 'id1', '-name', 'name1', '-value', 'value'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrCustom Property {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrCustom Property ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrCustom Property (['-interactive'])
```

setIdMgrLDAPAttrCache

The **setIdMgr LDAPA ttrCache** command configures the LDAP attribute cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-cachesDiskOffLoad

(String, optional)

-enabled

Indicates if you want to enable attribute caching. The default value is `true`. (Boolean, optional)

-cacheSize

The maximum size of the attribute cache defined by the number of attribute objects that are permitted in the attribute cache. The minimum value of this parameter is 100. The default value is 4000. (Integer, optional)

-cacheTimeout

The amount of time in seconds before the cached entries that are located in the attributes cache can be not valid. The minimum value of this parameter is 0. The attribute objects that are cached will remain in the attributes cache until the virtual member manager changes the attribute objects. The default value is 1200. (Integer, optional)

-attributeSizeLimit

An integer that represents the maximum number of attribute object values that can cache in the attributes cache.

Some attributes, for example, the member attribute, contain many values. The `attributeSizeLimit` parameter prevents the attributes cache to cache large attributes. The default value is 2000. (Integer, optional)

-serverTTLAttribute

The name of the `ttl` attribute that is supported by the LDAP server. The attributes cache uses the value of this attribute to determine when the cached entries in the attributes cache will time out.

The `ttl` attribute contains the time, in seconds, that any information from the entry should be kept by a client before it is considered stale and a new copy is fetched. A value of 0 implies that the object will not be cached. For more information about this attribute, go to: <http://www.ietf.org/proceedings/98aug/I-D/draft-ietf-asid-ldap-cache-01.txt>.

The `ttl` attribute is not supported by all LDAP servers. If this attribute is supported by an LDAP server, you can set the value of the `serverTTLAttribute` parameter to the name of the `ttl` attribute in order to allow the value of the `ttl` attribute to determine when cached entries will time out. The time out value for different entries in attributes cache can be different.

For example, if the value of the `serverTTLAttribute` parameter is `ttl` and the attributes cache retrieves attributes of a user from an LDAP server, it will also retrieve the value of the `ttl` attribute of this user. If the value is 200, the WMM uses this value to set the time out for the attributes of the user in the attributes cache instead of using the value of `cacheTimeout`. You can set different `ttl` attribute values for different users. (String, optional)

- Returns: None

Examples**Batch mode example usage:**

- Using Jacl:

```
$AdminTask setIdMgrLDAPAttr Cache {-id id1}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPAttr Cache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPAttr Cache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPAttr Cache {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPAttr Cache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPAttr Cache (['-interactive'])
```

setIdMgrLDAPContextPool

The **setIdMgr LDAPContextPool** command sets up the LDAP context pool configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-enabled

By default, the context pool is enabled. If you set this parameter to `false`, the context pool is disabled. When the context pool is disabled, new context instances will be created for each request. The default value is `true`. (Boolean, optional)

-initPoolSize

The number of context instances that the virtual member manager LDAP adapter creates when it creates the pool. The valid range for this parameter is 1 to 50. The default value is 1. (Integer, optional)

-maxPoolSize

The maximum number of context instances that the context pool will maintain. Context instances that are in use and those that are idle contribute to this number. When the pool size reaches this number, new context instances cannot be created for new requests. The new request is blocked until a context instance is released by another request or is removed. The request checks periodically if there are context instances available in the pool according to the amount of time that you specify using the `poolWaitTime` parameter.

The minimum value for this parameter is 0. There is no maximum value. Setting the value of this parameter to 0 means that there is no maximum size and a request for a pooled context instance will use an existing pooled idle context instance or a newly created pooled context instance. The default value is 20. (Integer, optional)

-prefPoolSize

The preferred number of context instances that the context pool will maintain. Context instances that are in use and those that are idle contribute to this number. When there is a request for the use of a pooled context instance and the pool size is less than the preferred size, the context pool creates and uses a new pooled context instance regardless of whether an idle connection is available. When a request finishes with a pooled context instance and the pool size is greater than the preferred size, the context pool closes and removes the pooled context instance from the pool.

The valid range for this parameter is from 0 to 100. Setting the value of this parameter to 0 means that there is no preferred size and a request for a pooled context instance results in a newly created context instance only if no idle ones are available. The default value is 3. (Integer, optional)

-poolTimeout

An integer that represents the number of milliseconds that an idle context instance may remain in the pool without being closed and removed from the pool. When a context instance is requested from the pool, if this context already exists in the pool for more than the time defined by `poolTimeout`, this connection will be closed no matter this context instance is stale or active. A new context instance will be created and put back to the pool after it has been released from the request.

The minimum value for this parameter is 0. There is no maximum value. Setting the value of this parameter to 0 means that the context instances in the pool will remain in the pool until they are staled. The context pool catches the communication exception and recreates a new context instance. The default value is 0. (Integer, optional)

-poolWaitTime

The time interval in milliseconds that the request waits until the context pool rechecks if there are idle

context instances available in the pool when the number of context instances reaches the maximum pool size. If no idle context instance, the request will continue waiting for the same period of time until next checking.

The minimum value for the `poolWaitout` parameter is 0. There is no maximum value. A value of 0 for this parameter means that the context pool will not check if idle context exists. The request will be notified when a context instance releases from other requests. The default value is 3000.(Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPCon textPool {-id id1}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPCon textPool ('[-id id1']')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPCon textPool (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPCon textPool {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPCon textPool ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPCon textPool (['-interactive'])
```

setIdMgrLDAPGroupConfig

The **setIdMgr LDAPGroupConfig** command sets up the LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-updateGroup Membership

Updates the group membership if the member is deleted or renamed. Some LDAP servers, for example, Domino server, do not clean up the membership of the user when a user is deleted or renamed. If you choose these LDAP server types in the `ldapServerType` property, the value of this parameter is set to `true`. Use this parameter to change the value. The default value is `false`. (Boolean, optional)

-name

The name of the membership attribute. For example, `memberOf` in an active directory server and `ibm-allGroups` in IDS. (String, optional)

-scope

The scope of the membership attribute. The following are the possible values for this parameter:

- `direct` - The membership attribute only contains direct groups. Direct groups contain the member and are not contained through a nested group. For example, if `group1` contains `group2`, `group2` contains `user1`, then `group2` is a direct group of `user1`, but `group1` is not a direct group of `user1`.
- `nested` - The membership attribute contains both direct groups and nested groups.
- `all` - The membership attribute contains direct groups, nested groups, and dynamic members.

The default value is `direct`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAP GroupConfig {-id id1}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAP GroupConfig ('[-id id1']')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAP GroupConfig (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPGroup Config {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPGroup Config ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPGroup Config (['-interactive'])
```

setIdMgrLDAPSearchResultCache

The **setIdMgr LDAPSearch ResultCache** command sets up the LDAP search result cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-cachesDiskOffLoad

Loads the attributes caches and the search results onto hard disk. By default, when the number of cache entries reaches the maximum size of the cache, cache entries are evicted to allow new entries to enter the caches. If you enable this parameter, the evicted cache entries will be copied to disk for future access. The default value is `false`. (Boolean, optional)

-enabled

Enables the search results cache. The default value is `true`. (Boolean, optional)

-cacheSize

The maximum size of the search results cache. The number of naming enumeration objects that can be put into the search results cache. The minimum value of this parameter is 100. The default value is 2000. (Integer, optional)

-cacheTimeOut

The amount of time in seconds before the cached entries in the search results cache can be not valid. The minimum value for this parameter is 0. A value of 0 means that the cached naming enumeration objects will stay in the search results cache until there are configuration changes. The default value is 600. (Integer, optional)

-searchResultSizeLimit

The maximum number of entries contained in the naming enumeration object that can be cached in the search results cache. For example, if the results from a search contains 2000 users, the search results will not cache in the search results cache if the value of the of this property is set to 1000. The default value is 1000. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPSearch ResultCache {-id id1}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPSearch ResultCache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPSearch ResultCache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrLDAPSearch ResultCache {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrLDAPSearch ResultCache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrLDAPSearch ResultCache (['-interactive'])
```

setIdMgrEntryMappingRepository

The **setIdMgr Entry Mapping Repository** command sets or updates an entry mapping repository configuration.

Parameters and return values

-dataSourceName

The name of the data source. The default value is jdbc/wimDS. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-databaseType

The type of the database. The default value is DB2. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-dbURL

The URL of the database. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-dbAdminId

The database administrator ID. (String, required if database type is not Apache Derby.)

-dbAdminPassword

The database administrator password. (String, required if database type is not Apache Derby.)

-JDBCClass

The JDBC driver class name. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrEntry MappingRepository {-dbAdminId dbid1 -dbAdminPassword pw1}
```

- Using Jython string:

```
AdminTask.setIdMgrEntry MappingRepository ('[-dbAdminId dbid1 -dbAdminPassword pw1]')
```

- Using Jython list:

```
AdminTask.setIdMgrEntry MappingRepository (['-dbAdminId', 'dbid1', '-dbAdmin Password', 'pw1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrEntryMapping Repository {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrEntryMapping Repository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrEntryMapping Repository (['-interactive'])
```

setIdMgrPropertyExtensionRepository

The **setIdMgr Property Extension Repository** command sets or updates the property extension repository configuration.

Parameters and return values

-dataSourceName

The name of the data source. The default value is jdbc/wimDS. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-databaseType

The type of the database. The default value is DB2. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-dbURL

The URL of the database. The parameter is required if the property extension is not set. The parameter is not required if the command is used to update the existing configuration. (String)

-dbAdminId

The database administrator ID. (String, required if database type is not Apache Derby.)

-dbAdminPassword

The database administrator password. (String, required if database type is not Apache Derby.)

-entityRetrievalLimit

The limit for the retrieval of entities. (Integer, required)

-JDBCDriverClass

The JDBC driver class name. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrProperty ExtensionRepository {-entity RetrievalLimit 10 -JDBC DriverClass classname}
```

- Using Jython string:

```
AdminTask.setIdMgrProperty ExtensionRepository ('[-entity RetrievalLimit 10 -JDBC DriverClass classname]')
```

- Using Jython list:

```
AdminTask.setIdMgrProperty ExtensionRepository (['-entity RetrievalLimit', '10', '-JDBC DriverClass', 'classname'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrProperty ExtensionRepository {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrProperty ExtensionRepository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.setIdMgrPropertyExt ensionRepository (['-interactive'])
```

updateIdMgrDBRepository

The **updateId MgrDB Repository** command updates the configuration for the database repository that you specify.

Parameters and return values

- id** The ID of the repository. (String, required)
- dataSourceName**
The name of the data source. The default value is jdbc/wimDS. (String, optional)
- databaseType**
The type of the database. The default value is DB2. (String, optional)
- dbURL**
The URL of the database. (String, optional)
- dbAdminId**
The database administrator ID. (String, optional)
- dbAdminPassword**
The database administrator password. (String, optional)
- entityRetrievalLimit**
Indicates the value of the retrieval limit on database entries. The default value is 200. (Integer, optional)
- JDBCDriverClass**
The JDBC driver class name. (String, optional)
- saltLength**
The salt length in bits. The default value is 12. (Integer, optional)
- encryptionKey**
The default value is rZ15ws0e1y9yHk3zCs3sTMv/ho8fY17s. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrDB Repository {-id id1}
```
- Using Jython string:

```
AdminTask.updateIdMgrDB Repository ('[-id id1']')
```
- Using Jython list:

```
AdminTask.updateIdMgrDB Repository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrDB Repository {-interactive}
```
- Using Jython string:

```
AdminTask.updateIdMgrDB Repository ('[-interactive]')
```
- Using Jython list:

```
AdminTask.updateIdMgrDB Repository (['-interactive'])
```

updateIdMgrFileRepository

The **updateId MgrFile Repository** command updates the configuration for the file repository that you specify. To update other properties of the file repository use the **update IdMgr Repository** command.

Parameters and return values

-id The ID of the repository. (String, required)

-messageDigest Algorithm

The message digest algorithm that will be used for hashing the password. The default value is SHA-1. Valid values include the following: SHA-1, SHA-384, or SHA-512.(String, optional)

-baseDirectory

The base directory where the file will be created in order to store the data. The default is to be dynamically built during run time using user.install.root and cell name. (String, optional)

-fileName

The file name of the repository. The default value is fileRegistry.xml. (String, optional)

-saltLength

The salt length of the randomly generated salt for password hashing. The default value is 12. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrFile Repository {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrFile Repository ('[-id id1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrFile Repository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrFile Repository {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrFile Repository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrFile Repository (['-interactive'])
```

updateIdMgrLDAPAttrCache

The **updateId MgrLDAP AttrCache** command updates the LDAP attribute cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-cachesDiskOffLoad

(String, optional)

-enabled

Indicates if you want to enable attribute caching. The default value is true. (Boolean, optional)

-cacheSize

The maximum size of the attribute cache defined by the number of attribute objects that are permitted in the attribute cache. The minimum value of this parameter is 100. The default value is 4000. (Integer, optional)

-cacheTimeOut

The amount of time in seconds before the cached entries that are located in the attributes cache can

be not valid. The minimum value of this parameter is 0. The attribute objects that are cached will remain in the attributes cache until the virtual member manager changes the attribute objects. The default value is 1200. (Integer, optional)

-attributeSizeLimit

An integer that represents the maximum number of attribute object values that can cache in the attributes cache.

Some attributes, for example, the member attribute, contain many values. The attributeSizeLimit parameter prevents the attributes cache to cache large attributes. The default value is 2000. (Integer, optional)

-serverTTLAttribute

The name of the ttl attribute that is supported by the LDAP server. The attributes cache uses the value of this attribute to determine when the cached entries in the attributes cache will time out.

The ttl attribute contains the time, in seconds, that any information from the entry should be kept by a client before it is considered stale and a new copy is fetched. A value of 0 implies that the object will not be cached. For more information about this attribute, go to: <http://www.ietf.org/proceedings/98aug/I-D/draft-ietf-asid-ldap-cache-01.txt>.

The ttl attribute is not supported by all LDAP servers. If this attribute is supported by an LDAP server, you can set the value of the serverTTLAttribute parameter to the name of the ttl attribute in order to allow the value of the ttl attribute to determine when cached entries will time out. The time out value for different entries in attributes cache can be different.

For example, if the value of the serverTTLAttribute parameter is ttl and the attributes cache retrieves attributes of a user from an LDAP server, it will also retrieve the value of the ttl attribute of this user. If the value is 200, the WMM uses this value to set the time out for the attributes of the user in the attributes cache instead of using the value of cacheTimeout. You can set different ttl attribute values for different users.

(String, optional)

- Returns: None

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP AttrCache {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP AttrCache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP AttrCache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP AttrCache {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP AttrCache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP AttrCache (['-interactive'])
```

updateIdMgrLDAPContextPool

The **updateId MgrLDAP ContextPool** command updates the LDAP context pool configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-enabled

By default, the context pool is enabled. If you set the value of this parameter to `false`, the context pool is disabled which means that a new context instance will be created for each request. The default value is `true`. (Boolean, optional)

-initPoolSize

The number of context instances that the virtual member manager LDAP adapter creates when it creates the pool. The valid range for this parameter is 1 to 50. The default value is 1. (Integer, optional)

-maxPoolSize

The maximum number of context instances that can be maintained concurrently by the context pool. Both in-use and idle context instances contribute to this number. When the pool size reaches this number, new context instances cannot be created for new requests. The new request is blocked until a context instance is released by another request or is removed. The request checks periodically if there are context instances available in the pool according to the value defined for the `poolWaitTime` parameter. The minimum value of the `maxPoolSize` parameter is 0. There is no maximum value. A maximum pool size of 0 means that there is no maximum size and that a request for a pooled context instance will use an existing pooled idle context instance or a newly created pooled context instance. The default value is 20. (Integer, optional)

-prefPoolSize

The preferred number of context instances that the Context Pool should maintain. Both in-use and idle context instances contribute to this number. When there is a request for the use of a pooled context instance and the pool size is less than the preferred size, Context Pool will create and use a new pooled context instance regardless of whether an idle connection is available. When a request is finished with a pooled context instance and the pool size is greater than the preferred size, the Context Pool will close and remove the pooled context instance from the pool. The valid range of the `prefPoolSize` parameter is 0 to 100. A preferred pool size of 0 means that there is no preferred size: A request for a pooled context instance will result in a newly created context instance only if no idle ones are available. The default value is 3. (Integer, optional)

-poolTimeout

An integer that represents the number of milliseconds that an idle context instance may remain in the pool without being closed and removed from the pool. When a context instance is requested from the pool, if this context already exists in the pool for more than the time defined by `poolTimeout`, this connection will be closed no matter if this context instance is stale or active. A new context instance will be created and put back to the pool after it has been released from the request. The minimum value of `poolTimeout` is 0. There is no maximum value. A `poolTimeout` of 0 means that the context instances in the pool will remain in the pool until they are staled. In this case, Context Pool will catch the communication exception and recreate a new context instance. The default value is 0. (Integer, optional)

-poolWaitTime

The time interval (in milliseconds) that the request will wait until the Context Pool checks again if there are idle context instances available in the pool when the number of context instances reaches the maximum pool size. If there is still no idle context instance, the request will continue waiting for the same period of time until next checking. The minimum value of `poolWaitTime` is 0. There is no maximum value. A `poolWaitTime` of 0 means the Context Pool will not check if there are idle context instances. Instead, the request will be notified when there is a context instance released from other requests. The default value is 3000. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP ContextPool {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP ContextPool ('[-id id1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP ContextPool (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP ContextPool {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP ContextPool ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP ContextPool (['-interactive'])
```

updateIdMgrLDAPEntityType

The **updateId MgrLDAP EntityType** command updates an existing LDAP entity type definition to LDAP repository configuration. You can use this command to add more values to multi-valued parameters. If the property already exists, the value of the property will be replaced. If the property does not exist, it will be added.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the entity type. (String, required)

-searchFilter

The search filter that you want to use to search the entity type. (String, optional)

-objectClasses

One or more object classes for the entity type. (String, optional)

-objectClassesForCreate

The object class that will be when you create an entity type object. You do not have to specify the value of this parameter if it is the same as the value of the objectClasses parameter. (String, optional)

-searchBases

The search base or bases to use while searching the entity type. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAPEntityType Type {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPEntityType Type ('[-id id1 -name name1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPEntityType Type (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP EntityType {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP EntityType ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP EntityType (['-interactive'])
```

updateIdMgrLDAPGroupDynamicMemberAttr

The **updateIdMgr LDAPGroup Dynamic MemberAttr** command updates a dynamic member attribute configuration to an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the LDAP attribute that is used as the group member attribute. For example, memberURL. (String, required)

-objectClass

The group object class that contains the dynamic member attribute. For example groupOfURLs. If you do not define this parameter, the dynamic member attribute will apply to all group object classes. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAPGroup DynamicMemberAttr {-id id1 -name name1 -objectClass groupOfURLs}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPGroup DynamicMemberAttr ('[-id id1 -name name1 -objectClass groupOfURLs]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPGroup DynamicMemberAttr (['-id', 'id1', '-name', 'name1', '-objectClass', 'groupOfURLs'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAPGroup DynamicMemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPGroup DynamicMemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPGroup DynamicMemberAttr (['-interactive'])
```

updateIdMgrLDAPGroupMemberAttr

The **updateIdMgr LDAPGroup MemberAttr** command updates a member attribute configuration of an LDAP group configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The name of the LDAP attribute that is used as the group member attribute. For example, member or uniqueMember. (String, required)

-objectClass

The group object class that contains the member attribute. For example, `groupOfNames` or `groupOfUniqueNames`. If you do not define this parameter, the member attribute applies to all group object classes. (String, optional)

-scope

The scope of the member attribute. The following are the valid values:

- `direct` - The member attribute only contains direct members whereby the member is directly contained by the group and not contained in a nested group. For example, if `group1` contains `group2`, `group2` contains `user1`, then `group2` is a direct member of `group1` but `user1` is not a direct member of `group1`. Both `member` and `uniqueMember` are direct member attributes.
- `nested` - The member attribute contains both direct members and nested members.

-dummyMember

When you create a group without specifying a member, a dummy member will be filled in automatically to avoid receiving an exception that indicates that there is a mandatory attribute missing. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP GroupMemberAttr {-id id1 -name name1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP GroupMemberAttr ('[-id id1 -name name1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP GroupMemberAttr (['-id', 'id1', '-name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAPGroup MemberAttr {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPGroup MemberAttr ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPGroup MemberAttr (['-interactive'])
```

updateIdMgrLDAPRepository

The **updateId MgrLDAP Repository** command updates an LDAP repository configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-ldapServerType

The type of LDAP server that is being used. The default value is `IDS51`. (String, optional)

-adapterClassName

The default value is `com.ibm.ws.wim.adapter.ldap.LdapAdapter`. (String, optional)

-certificateMapMode

Specifies whether to map X.509 certificates into a LDAP directory by exact distinguished name or by certificate filter. The default value is `exactdn`. To use the certificate filter for the mapping, specify `certificatefilter`. (String, optional)

-certificateFilter

If `certificateMapMode` has the value `certificatefilter`, then this property specifies the LDAP filter which maps attributes in the client certificate to entries in LDAP. (String, optional)

-isExtIdUnique

Specifies if the external ID is unique. The default value is `true`. (Boolean, optional)

-loginProperties

Indicates the property name used for login. (String, optional)

-primaryServerQueryTimeInterval

Indicates the polling interval for testing the primary server availability. The value of this parameter is specified in minutes. The default value is 15. (Integer, optional)

-returnToPrimaryServer

Indicates to return to the primary LDAP server when it is available. The default value is `true`. (Boolean, optional)

-supportAsyncMode

Indicates if the async mode is supported or not. The default value is `false`. (Boolean, optional)

-supportSorting

Indicates if sorting is supported or not. The default value is `false`. (Boolean, optional)

-supportPaging

Indicates if paging is supported or not. The default value is `false`. (Boolean, optional)

-supportTransactions

Indicates if transactions are supported or not. The default value is `false`. (Boolean, optional)

-supportExternalName

Indicates if external names are supported or not. The default value is `false`. (Boolean, optional)

-sslConfiguration

The SSL configuration. (String, optional)

-translateRDN

Indicates to translate RDN or not. The default value is `false`. (Boolean, optional)

-searchTimeLimit

The value of search time limit. (Integer, optional)

-searchCountLimit

The value of search count limit. (Integer, optional)

-searchPageSize

The value of search page size. (Integer, optional)

Examples**Batch mode example usage:**

- Using Jacl:

```
$AdminTask updateIdMgrLDAP Repository {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP Repository ('[-id id1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP Repository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP Repository {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP Repository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP Repository (['-interactive'])
```

updateIdMgrLDAPSearchResultCache

The **updateIdMgr LDAPSearch ResultCache** command updates the LDAP search result cache configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-cachesDiskOffLoad

Loads the attributes caches and the search results onto hard disk. By default, when the number of cache entries reaches the maximum size of the cache, cache entries are evicted to allow new entries to enter the caches. If you enable this parameter, the evicted cache entries will be copied to disk for future access. The default value is `false`. (Boolean, optional)

-enabled

Enables the search results cache. The default value is `true`. (Boolean, optional)

-cacheSize

The maximum size of the search results cache. The number of naming enumeration objects that can be put into the search results cache. The minimum value of this parameter is 100. The default value is 2000. (Integer, optional)

-cacheTimeOut

The amount of time in seconds before the cached entries in the search results cache can be not valid. The minimum value for this parameter is 0. A value of 0 means that the cached naming enumeration objects will stay in the search results cache until there are configuration changes. The default value is 600. (Integer, optional)

-searchResultSizeLimit

The maximum number of entries contained in the naming enumeration object that can be cached in the search results cache. For example, if the results from a search contains 2000 users, the search results will not cache in the search results cache if the value of the of this property is set to 1000. The default value is 1000. (Integer, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP SearchResultCache {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPSearch ResultCache ('[-id id1']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPSearch ResultCache (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP SearchResultCache {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPSearch ResultCache ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPSearch ResultCache (['-interactive'])
```

updateIdMgrLDAPServer

The **updateIdMgr LDAPServer** command updates an LDAP server configuration for the LDAP repository ID that you specify.

Parameters and return values

-id The ID of the repository. (String, required)

-host

The host name for the LDAP server that contains the properties that you want to modify. (String, required)

-port

The port number for the LDAP server. (Integer, optional)

-authentication

Indicates the authentication method to use. The default value is `simple`. Valid values include: `none` or `strong`. (String, optional)

-bindDN

The binding domain name for the LDAP server. (String, optional)

-bindPassword

The binding password. The password is encrypted before it is stored. (String, optional)

-certificateMapMode

Specifies whether to map X.509 certificates into a LDAP directory by exact distinguished name or by certificate filter. The default value is `exactdn`. To use the certificate filter for the mapping, specify `certificatefilter`. (String, optional)

-certificateFilter

If `certificateMapMode` has the value `certificatefilter`, then this property specifies the LDAP filter which maps attributes in the client certificate to entries in LDAP. (String, optional)

-connectTimeout

The connection timeout measured in seconds. The default value is 0. (Integer, optional)

-connectionPool

The connection pool. The default value is `false`. (Boolean, optional)

-derefAliases

Controls how aliases are dereferenced. The default value is `always`. Valid values include:

- `never` - never deference aliases
- `finding` - deferences aliases only during name resolution
- `searching` - deferences aliases only after name resolution

(String, optional)

-ldapServerType

The type of LDAP server being used. The default value is `IDS51`. (String, optional)

-primary_host

The host name for the primary LDAP server. (String, optional)

-referral

The LDAP referral. The default value is `ignore`. Valid values include: `follow`, `throw`, or `false`. (String, optional)

-sslConfiguration

The SSL configuration. (String, optional)

-sslEnabled

Indicates to enable SSL or not. The default value is `false`. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAPServer {-id id1 -host myhost.ibm.com}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAPServer ('[-id id1 -host myhost.ibm.com']')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAPServer (['-id', 'id1', '-host', 'myhost.ibm.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrLDAP Server {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrLDAP Server ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrLDAP Server (['-interactive'])
```

updateIdMgrRepository

The **updateIdMgr Repository** command updates the common repository configuration.

Parameters and return values

-id The ID of the repository. (String, required)

-adapterClassName

The implementation class name for the repository adapter. (String, optional)

-EntityTypesNot AllowCreate

The name of the entity type that should not be created in this repository. (String, optional)

-EntityTypesNotAllowUpdate

The name of the entity type that should not be updated in this repository. (String, optional)

-EntityTypesNotAllowRead

The name of the entity type that should not be read from this repository. (String, optional)

-EntityTypesNotAllowDelete

The name of the entity type that should not be deleted from this repository. (String, optional)

-loginProperties

(String, optional)

-readOnly

Indicates if this is a read only repository. The default value is false. (Boolean, optional)

-repositoriesForGroups

The repository ID where group data is stored. (String, optional)

-supportPaging

Indicates if the repository supports paging or not. (Boolean, optional)

-supportSorting

Indicates if the repository supports sorting or not. (Boolean, optional)

-supportTransactions

Indicates if the repository supports transaction or not. (Boolean, optional)

-isExtIdUnique

Specifies if the external ID is unique or not. (Boolean, optional)

-supportedExternalName

Indicates if the repository supports external names or not. (Boolean, optional)

-supportAsyncMode

Indicates if the adapter supports async mode or not. The default value is false. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrRepository {-id id1}
```

- Using Jython string:

```
AdminTask.updateIdMgrRepository ('[-id id1]')
```

- Using Jython list:

```
AdminTask.updateIdMgrRepository (['-id', 'id1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrRepository {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrRepository ('[-interactive]')
```

- Using Jython list:

```
AdminTask.updateIdMgrRepository (['-interactive'])
```

updateIdMgrRepositoryBaseEntry

The **updateIdMgr Repository BaseEntry** command updates a base entry to the specified repository.

Parameters and return values

-id The ID of the repository. (String, required)

-name

The distinguished name of a base entry. (String, required)

-nameInRepository

The distinguished name in the repository that uniquely identifies the base entry name. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrRepositoryBaseEntry {-id id1 name name1}
```

- Using Jython string:

```
AdminTask.updateIdMgrRepositoryBaseEntry ('[-id id1 name name1]')
```

- Using Jython list:

```
AdminTask.updateIdMgrRepositoryBaseEntry (['-id', 'id1', 'name', 'name1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask updateIdMgrRepositoryBaseEntry {-interactive}
```

- Using Jython string:

```
AdminTask.updateIdMgrRepositoryBaseEntry ('[-interactive]')
```


- Using Jython list:

```
AdminTask.updateIdMgrRepositoryBaseEntry (['-interactive'])
```

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

IdMgrRealmConfig command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure the member manager realm and realms. The commands and parameters in the IdMgrRealmConfig group can be used to create and manage your realm configuration.

The IdMgrRealmConfig command group for the AdminTask object includes the following commands:

- “addIdMgrRealmBaseEntry”
- “createIdMgrRealm” on page 1211
- “deleteIdMgrRealm” on page 1212
- “deleteIdMgrRealmBaseEntry” on page 1212
- “getIdMgrDefaultRealm” on page 1213
- “getIdMgrRepositoriesForRealm” on page 1214
- “getIdMgrRealm” on page 1214
- “listIdMgrRealms” on page 1215
- “listIdMgrRealmBaseEntries” on page 1215
- “renameIdMgrRealm” on page 1216
- “setIdMgrDefaultRealm” on page 1217
- “updateIdMgrRealm” on page 1217

addIdMgrRealmBaseEntry

The **addIdMgrRealmBaseEntry** command adds a base entry to a specific realm configuration and links the realm with the repository.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

-baseEntry

Specifies the name of the base entry. (String, optional)

Optional parameters

None

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrRealmBaseEntry {-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository}
```
- Using Jython string:

```
AdminTask.addIdMgrRealmBaseEntry ('[-name defaultWIMFileBasedRealm -baseEntry o=sampleFileRepository]')
```
- Using Jython list:

```
AdminTask.addIdMgrRealmBaseEntry (['-name', 'defaultWIMFileBasedRealm', '-baseEntry', 'o=sampleFileRepository'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addIdMgrRealmBaseEntry {-interactive}
```
- Using Jython string:

```
AdminTask.addIdMgrRealmBaseEntry ('[-interactive]')
```

createIdMgrRealm

The **createIdMgrRealm** command creates a realm configuration.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Optional parameters

-securityUse

Specifies a string that indicates if this virtual realm will be used in security now, later, or never. The default value is active. Additional values includes: inactive and nonSelectable. (String, optional)

-delimiter

Specifies the delimiter used for this realm. The default value is /. (String, optional)

-allowOperationIfReposDown

Specifies whether the system allows a repository operation such as get or search to complete successfully, even if repositories in the realm are down. The default value is false. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrRealm {-name realm1 -allowOperationIfReposDown true}
```
- Using Jython string:

```
AdminTask.createIdMgrRealm ('[-name realm1 -allowOperationIfReposDown true]')
```
- Using Jython list:

```
AdminTask.createIdMgrRealm (['-name', 'realm1', '-allowOperationIfReposDown', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createIdMgrRealm {-interactive}
```
- Using Jython string:

```
AdminTask.createIdMgrRealm ('[-interactive]')
```

deleteIdMgrRealm

The **deleteIdMgrRealm** command deletes the realm configuration that you specified.

Target Object

None.

Required parameters

-name

The realm name. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrRealm {-name realm1}
```
- Using Jython string:

```
AdminTask.deleteIdMgrRealm ('[-name realm1]')
```
- Using Jython list:

```
AdminTask.deleteIdMgrRealm (['-name', 'realm1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteIdMgrRealm {-interactive}
```
- Using Jython string:

```
AdminTask.deleteIdMgrRealm ('[-interactive]')
```
- Using Jython list:

```
AdminTask.deleteIdMgrRealm (['-interactive'])
```

deleteIdMgrRealmBaseEntry

The **deleteIdMgrRealmBaseEntry** command deletes a base entry from a realm configuration that you specified.

The realm must always contain at least one base entry, thus you cannot remove every entry.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

-baseEntry

Specifies the name of a base entry. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteIdMgrRealmBaseEntry {-name realm1 -baseEntry entry1}`
- Using Jython string:
`AdminTask.deleteIdMgrRealmBaseEntry ('[-name realm1 -baseEntry entry1]')`
- Using Jython list:
`AdminTask.deleteIdMgrRealmBaseEntry (['-name', 'realm1', '-baseEntry', 'entry1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteIdMgrRealmBaseEntry {-interactive}`
- Using Jython string:
`AdminTask.deleteIdMgrRealmBaseEntry ('[-interactive]')`

getIdMgrDefaultRealm

The **getIdMgrDefaultRealm** command returns the default realm name.

Target Object

None.

Required parameters

None.

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask getIdMgrDefaultRealm`
- Using Jython string:
`AdminTask.getIdMgrDefaultRealm()`
- Using Jython list:
`AdminTask.getIdMgrDefaultRealm()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask getIdMgrDefaultRealm {-interactive}`
- Using Jython string:
`AdminTask.getIdMgrDefaultRealm ('[-interactive]')`

getIdMgrRepositoriesForRealm

The **getIdMgrRepositoriesForRealm** command returns repository specific details for the repositories configured for a specified realm.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask getIdMgrRepositoriesForRealm {-name realm1}`
- Using Jython string:
`AdminTask.getIdMgrRepositoriesForRealm ('[-name realm1']')`
- Using Jython list:
`AdminTask.getIdMgrRepositoriesForRealm (['-name', 'realm1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask getIdMgrRepositoriesForRealm {-interactive}`
- Using Jython string:
`AdminTask.getIdMgrRepositoriesForRealm ('[-interactive]')`

getIdMgrRealm

The **getIdMgrRealm** command returns the configuration parameters for the realm that you specified.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrRealm {-name realm1}
```

- Using Jython string:

```
AdminTask.getIdMgrRealm ('[-name realm1]')
```

- Using Jython list:

```
AdminTask.getIdMgrRealm (['-name', 'realm1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getIdMgrRealm {-interactive}
```

- Using Jython string:

```
AdminTask.getIdMgrRealm ('[-interactive]')
```

listIdMgrRealms

The **listIdMgrRealms** command returns all of the names of the configured realms.

Target Object

None.

Required parameters

None.

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrRealms
```

- Using Jython string:

```
AdminTask.listIdMgrRealms()
```

- Using Jython list:

```
AdminTask.listIdMgrRealms()
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listIdMgrRealms {-interactive}
```

- Using Jython string:

```
AdminTask.listIdMgrRealms ('[-interactive]')
```

listIdMgrRealmBaseEntries

The **listIdMgrRealmBaseEntries** command returns all of the names of the configured realms.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listIdMgrRealmBaseEntries {-name realm1}`
- Using Jython string:
`AdminTask.listIdMgrRealmBaseEntries ('[-name realm1']')`
- Using Jython list:
`AdminTask.listIdMgrRealmBaseEntries (['-name', 'realm1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listIdMgrRealmBaseEntries {-interactive}`
- Using Jython string:
`AdminTask.listIdMgrRealmBaseEntries ('[-interactive]')`

renameIdMgrRealm

The **renameIdMgrRealm** command renames the name of the realm that you specified.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Optional parameters

None.

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask renameIdMgrRealm {-name realm1}`
- Using Jython string:
`AdminTask.renameIdMgrRealm ('[-name realm1']')`
- Using Jython list:
`AdminTask.renameIdMgrRealm (['-name', 'realm1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask renameIdMgrRealm {-interactive}`
- Using Jython string:
`AdminTask.renameIdMgrRealm ('[-interactive]')`

- Using Jython list:

```
AdminTask.renameIdMgrRealm (['-interactive'])
```

setIdMgrDefaultRealm

The **setIdMgrDefaultRealm** command sets up the default realm configuration.

Parameters and return values

-name

specifies the name of the realm that is used as a default realm when the caller does not specify any in context. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrDefaultRealm {-name realm1}
```

- Using Jython string:

```
AdminTask.setIdMgrDefaultRealm (['-name realm1'])
```

- Using Jython list:

```
AdminTask.setIdMgrDefaultRealm (['-name', 'realm1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask setIdMgrDefaultRealm {-interactive}
```

- Using Jython string:

```
AdminTask.setIdMgrDefaultRealm (['-interactive'])
```

updateIdMgrRealm

The **updateIdMgrRealm** command updates the configuration for a realm that you specify.

Target Object

None.

Required parameters

-name

Specifies the name of the realm. (String, required)

Optional parameters

-securityUse

Specifies a string that indicates if this realm will be used in security now, later, or never. The default value is `active`. Additional values includes: `inactive` and `nonSelectable`. (String, optional)

-delimiter

specifies the delimiter used for this realm. The default value is `/`. (String, optional)

-allowOperationIfReposDown

Specifies whether the system allows a repository operation such as `get` or `search` to complete successfully, even if repositories in the realm are down. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask updateIdMgrRealm {-name realm1}`
- Using Jython string:
`AdminTask.updateIdMgrRealm ('[-name realm1]')`
- Using Jython list:
`AdminTask.updateIdMgrRealm (['-name', 'realm1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask updateIdMgrRealm {-interactive}`
- Using Jython string:
`AdminTask.updateIdMgrRealm ('[-interactive]')`
- Using Jython list:
`AdminTask.updateIdMgrRealm (['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

“IdMgrRepositoryConfig command group for the AdminTask object” on page 1163

You can use the Jython or Jacl scripting languages to configure security. The commands and parameters in the IdMgrRepositoryConfig group can be used to create and manage the virtual member manager and LDAP directory properties.

“IdMgrConfig command group for the AdminTask object” on page 1158

You can use the Jython or Jacl scripting languages to configure the virtual member manager with the wsadmin tool. The commands and parameters in the IdMgrConfig group can be used to create and manage your entity type configuration.

WIMManagementCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the WIMManagementCommands group can be used to create and manage groups, members, and users in the virtual member manager.

The WIMManagementCommands command group for the AdminTask object includes the following commands:

- “createGroup” on page 1219
- “createUser” on page 1220
- “deleteGroup” on page 1221
- “deleteUser” on page 1221
- “duplicateMembershipOfGroup” on page 1222
- “duplicateMembershipOfUser” on page 1222
- “getGroup” on page 1223
- “getMembershipOfGroup” on page 1224
- “getMembershipOfUser” on page 1224
- “getMembersOfGroup” on page 1225

- “getUser” on page 1225
- “removeMemberFromGroup” on page 1226
- “searchGroups” on page 1226
- “searchUsers” on page 1227
- “updateGroup” on page 1228
- “updateUser” on page 1229

createGroup

The **createGroup** command creates a new group in the virtual member manager. After the command completes, the new group will appear in the repository. For LDAP, a group must contain a member. The `memberUniqueName` parameter is optional in this case. If you set the `memberUniqueName` parameter to the unique name of a group or a user, the group or user will be added as a member of the group.

Parameters and return values

-cn

Specifies the common name for the group that you want to create. This parameter maps to the `cn` property in virtual member manager. (String, required)

-description

Specifies additional information about the group that you want to create. This parameter maps to the `description` property in a virtual member manager object. (String, optional)

-parent

Specifies the repository in which you want to create the group. This parameter maps to the `parent` property in the virtual member manager. (String, optional)

-memberUniqueName

Specifies the unique name value for the user or group that you want to add to the new group. This parameter maps to the `uniqueName` property in the virtual member manager. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask createGroup {-cn groupA -description a group of admins}
```
- Using Jython string:


```
AdminTask.createGroup ('[-cn groupA -description a group of admins]')
```
- Using Jython list:


```
AdminTask.createGroup (['-cn', 'groupA', '-description', 'a group of admins'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask createGroup {-interactive}
```
- Using Jython string:


```
AdminTask.createGroup ('[-interactive]')
```
- Using Jython list:


```
AdminTask.createGroup (['-interactive'])
```

createUser

The **createUser** command creates a new user in the default repository or a repository that the parent command parameter specifies. This command creates a person entity and a login account entity in the virtual member manager.

Parameters and return values

-uid

Specifies the unique ID for the user that you want to create. Virtual member manager then creates a `uniqueId` value and a `uniqueName` value for the user. This parameter maps to the `uid` property in the virtual member manager. (String, required)

-password

Specifies the password for the user. This parameter maps to the `password` property in the virtual member manager. (String, required)

-confirmPassword

Specifies the password again to validate how it was entered for the `password` parameter. This parameter maps to the `password` property in virtual member manager. (String, optional)

-cn

Specifies the first name or given name of the user. This parameter maps to the `cn` property in virtual member manager. (String, optional)

-surname

Specifies the last name or family name of the user. This parameter maps to the `sn` property in virtual member manager. (String, optional)

-ibm-primaryEmail

Specifies the e-mail address of the user. This parameter maps to the `ibm-PrimaryEmail` property in the virtual member manager. (String, optional)

-parent

Specifies the repository in which you want to create the user. This parameter maps to the `parent` property in the virtual member manager. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createUser {-uid 123 -password tempPass -confirmPassword tempPass -cn Jane -sn Doe -mail janedoe@ acme.com}
```
- Using Jython string:

```
AdminTask.createUser ('[-uid 123 -password tempPass -confirmPassword tempPass -cn Jane -sn Doe -mail janedoe@ acme.com]')
```
- Using Jython list:

```
AdminTask.createUser (['-uid', '123', '-password', 'tempPass', '-confirmPassword', 'tempPass', '-cn', 'Jane', '-sn', 'Doe', '-ibm -mail', 'janedoe@ acme.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createUser {-interactive}
```
- Using Jython string:

```
AdminTask.createUser ('[-interactive]')
```
- Using Jython list:

```
AdminTask.createUser (['-interactive'])
```

deleteGroup

The **deleteGroup** command deletes a group in the virtual member manager. You cannot use this command to delete descendants. When this command completes, the group will be deleted from the repository.

Parameters and return values

-uniqueName

Specifies the unique name value for the group that you want to delete. This parameter maps to the `uniqueName` property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteGroup {-uniqueName cn=opera tors,cn=users,dc=yourco, dc=com}
```
- Using Jython string:

```
AdminTask.deleteGroup ('[-uniqueName cn=ope rators,cn=users,dc=you rco,dc=com]')
```
- Using Jython list:

```
AdminTask.deleteGroup (['-uniqueName', 'cn =operators,cn=users,dc =yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteGroup {-interactive}
```
- Using Jython string:

```
AdminTask.deleteGroup ('[-interactive]')
```
- Using Jython list:

```
AdminTask.deleteGroup (['-interactive'])
```

deleteUser

The **deleteUser** command deletes a user from the virtual member manager. This includes a person object and an account object in the non-merged repositories.

Parameters and return values

-uniqueName

Specifies the unique name value for the user that you want to delete. This parameter maps to the `uniqueName` property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteUser {-uniqueName uid=dmey ers,cn=users,dc=yourco, dc=com}
```
- Using Jython string:

```
AdminTask.deleteUser ('[-uniqueName uid= dmeyers,cn=users,dc= yourco,dc=com]')
```
- Using Jython list:

```
AdminTask.deleteUser ([' -uniqueName', 'uid=dm eyers,cn=users,dc=yourco, dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteUser {-interactive}
```

- Using Jython string:

```
AdminTask.deleteUser ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteUser (['-interactive'])
```

duplicateMembershipOfGroup

Use the **duplicate Membership OfGroup** command to make a one group a member of all of the same groups as another group. For example, group A is in group B and group C. To add group D to the same groups as group A, use the **duplicate Membership OfGroup** command.

Parameters and return values

-copyToName

Specifies the name of the group to which you want to add the memberships of the group specified in the copyFromName parameter. (String, required)

-copyFromName

Specifies the name of the group from which you want to copy the group memberships for another group to use. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask duplicateMember shipOfGroup {-copyToName cn=operators,cn=groups, dc=yourco,dc=com  
-copy FromName cn=admins,cn= groups,dc=yourco,dc=com}
```

- Using Jython string:

```
AdminTask.duplicateMembers hipOfGroup ('[-copyToName cn=operators,cn=groups, dc=yourco,dc=com  
-copy FromName cn=admins,cn=gr oups,dc=yourco,dc=com]')
```

- Using Jython list:

```
AdminTask.duplicateMember shipOfGroup (['-copyToName', 'cn=operators,cn=groups, dc=yourco,dc=com',  
'-copy FromName', 'cn=admins,cn =groups,dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask duplicateMember shipOfGroup {-interactive}
```

- Using Jython string:

```
AdminTask.duplicateMember shipOfGroup ('[-interactive]')
```

- Using Jython list:

```
AdminTask.duplicateMember shipOfGroup (['-interactive'])
```

duplicateMembershipOfUser

Use the **duplicate Membership OfUser** command to make a one user a member of all of the same groups as another user. For example, user 1 is in group B and group C. To add user 2 to the same groups as user 1, use the **duplicate Membership OfUser** command.

Parameters and return values

-copyToName

Specifies the name of the user to which you want to add the memberships of the user specified in the copyFromName parameter. (String, required)

-copyFromName

Specifies the name of the user from which you want to copy the group memberships for another user to use. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask duplicateMember shipOfUser {-copyToName uid=meyersd,cn=users,dc=yourco,dc=com  
-copy FromName uid=jhart,cn=users,dc=yourco,dc=com}
```

- Using Jython string:

```
AdminTask.duplicateMember shipOfUser (['-copyToName uid=meyersd,cn=users,dc=yourco,dc=com',  
-copy FromName uid=jhart,cn=users,dc=yourco,dc=com'])
```

- Using Jython list:

```
AdminTask.duplicateMember shipOfUser (['-copyToName', 'uid=meyersd,cn=users,dc=yourco,dc=com',  
'-copyFromName', 'uid=jhart,cn=users,dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask duplicateMember shipOfUser {-interactive}
```

- Using Jython string:

```
AdminTask.duplicateMember shipOfUser (['-interactive'])
```

- Using Jython list:

```
AdminTask.duplicateMember shipOfUser (['-interactive'])
```

getGroup

The **getGroup** command retrieves the common name and description of a group.

Parameters and return values

-uniqueName

Specifies the unique name value for the group that you want to view. This parameter maps to the `uniqueName` property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getGroup {-uniqueName cn=operators, cn=groups,dc=yourco,dc=com}
```

- Using Jython string:

```
AdminTask.getGroup (['-uniqueName cn=operators, cn=groups,dc=yourco, dc=com'])
```

- Using Jython list:

```
AdminTask.getGroup (['-uniqueName', 'cn=operators,cn=groups,dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getGroup {-interactive}
```

- Using Jython string:

```
AdminTask.getGroup (['-interactive'])
```

- Using Jython list:

```
AdminTask.getGroup (['-interactive'])
```

getMembershipOfGroup

The **getMembership OfGroup** command retrieves the groups of which a group is a member.

Parameters and return values

-uniqueName

Specifies the unique name value for the group whose group memberships you want to view. This parameter maps to the uniqueName property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getMembership OfGroup {-uniqueName uid=dmeyers,cn=users, dc=yourco,dc=com}
```
- Using Jython string:

```
AdminTask.getMembership OfGroup ('[-uniqueName uid=dmeyers,cn=users, dc=yourco,dc=com]')
```
- Using Jython list:

```
AdminTask.getMembership OfGroup (['-uniqueName', 'uid=dmeyers,cn=users, dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getMembership OfGroup {-interactive}
```
- Using Jython string:

```
AdminTask.getMembership OfGroup ('[-interactive]')
```
- Using Jython list:

```
AdminTask.getMembership OfGroup (['-interactive'])
```

getMembershipOfUser

The **getMembership OfUser** command retrieves the groups of which a user is a member.

Parameters and return values

-uniqueName

Specifies the unique name value for the user whose group memberships you want to view. This parameter maps to the uniqueName property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getMembership OfUser {-uniqueName uid=dmeyers,cn=users, dc=yourco,dc=com}
```
- Using Jython string:

```
AdminTask.getMembership OfUser ('[-uniqueName uid=dmeyers,cn=users, dc=yourco,dc=com]')
```
- Using Jython list:

```
AdminTask.getMembership OfUser (['-uniqueName', 'uid=dmeyers,cn=users, dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getMembership OfUser {-interactive}
```
- Using Jython string:

```
AdminTask.getMembership OfUser ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getMembership OfUser (['-interactive'])
```

getMembersOfGroup

The **getMembers OfGroup** command retrieves the members of a group.

Parameters and return values

-uniqueName

Specifies the unique name value for the group whose members you want to view. This parameter maps to the uniqueName property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getMembersOf Group {-uniqueName cn=operators,cn=groups ,dc=yourco,dc=com}
```

- Using Jython string:

```
AdminTask.getMembersOf Group ['(-uniqueName cn=operators,cn=groups ,dc=yourco,dc=com)']
```

- Using Jython list:

```
AdminTask.getMembersOf Group [('-uniqueName', 'cn=operators,cn=groups ,dc=yourco,dc=com')]
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask getMembersOfGroup {-interactive}
```

- Using Jython string:

```
AdminTask.getMembersOfGroup ('[-interactive]')
```

- Using Jython list:

```
AdminTask.getMembersOfGroup (['-interactive'])
```

getUser

The **getUser** command retrieves information about a user in the virtual member manager.

Parameters and return values

-uniqueName

Specifies the unique name value for the user that you want to view. This parameter maps to the uniqueName property in the virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask getUser {-user Name uid=dmeyers,cn=users,dc=yourco,dc=com}
```

- Using Jython string:

```
AdminTask.getUser ('[-use rName uid=dmeyers,cn= users,dc=yourco,dc=com]')
```

- Using Jython list:

```
AdminTask.getUser (['-use rName', 'uid=dmeyers,c n=users,dc=yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask getUser {-interactive}
```
- Using Jython string:


```
AdminTask.getUser ('[-interactive]')
```
- Using Jython list:


```
AdminTask.getUser (['-interactive'])
```

removeMemberFromGroup

The **removeMember FromGroup** command removes a user or a group from a group.

Parameters and return values

-memberUniqueName

Specifies the unique name value for the user or group that you want to remove from the specified group. This parameter maps to the uniqueName property in virtual member manager. (String, required)

-groupUniqueName

Specifies the unique name value for the group from which you want to remove the user or group that you specified with the memberUniqueName paramter. This parameter maps to the uniqueName property in virtual member manager. (String, required)

Examples

Batch mode example usage:

- Using Jacl:


```
$AdminTask removeMember FromGroup {-memberUnique Name uid=meyersd,cn= users,dc=yourco,dc=com  
-groupUniqueName cn= admins,cn-groups,dc= yourco,dc=com}
```
- Using Jython string:


```
AdminTask.removeMemberF romGroup ('[-memberUnique Name uid=meyersd,cn= users,dc=yourco,dc=com  
-groupUniqueName cn=a dmins,cn-groups,dc=your co,dc=com]')
```
- Using Jython list:


```
AdminTask.removeMemberFrom Group (['-memberUniqueName', 'uid=meyersd,cn=users, dc=yourco,dc=com',  
'-groupUniqueName', 'cn=admins,cn-groups,dc= yourco,dc=com'])
```

Interactive mode example usage:

- Using Jacl:


```
$AdminTask removeMember FromGroup {-interactive}
```
- Using Jython string:


```
AdminTask.removeMemberFr omGroup ('[-interactive]')
```
- Using Jython list:


```
AdminTask.removeMemberFr omGroup (['-interactive'])
```

searchGroups

Use the **searchGroups** command to find groups in the virtual member manager that match criteria that you provide. For example, you can use the **searchGroups** command to find all of the groups with a common name that begins with IBM. You can search for any virtual member manager property because the command is generic.

Parameters and return values

-cn

The first name or given name of the user. This parameter maps to the cn property in the virtual member manager. You must set this parameter or the description parameter, but not both. (String, optional)

-description

Specifies information about the group. This parameter maps to the description entity in a virtual member manager object. You must set this parameter or the cn parameter, but not both. (String, optional)

-timeLimit

Specifies the maximum amount of time in milliseconds that the search can run. The default value is no time limit. (String, optional)

-countLimit

Specifies the maximum number of results that you want returned from the search. By default, all groups found in the search are returned. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask searchGroups {cn *IBM*}`
- Using Jython string:
`AdminTask.searchGroups ('[cn *IBM*]')`
- Using Jython list:
`AdminTask.searchGroups (['cn', '*IBM*'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask searchGroups {-interactive}`
- Using Jython string:
`AdminTask.searchGroups (['-interactive'])`
- Using Jython list:
`AdminTask.searchGroups (['-interactive'])`

searchUsers

Use the **searchUsers** command to find users in the virtual member manager that match criteria that you provide. For example, you can use the **searchUsers** command to find all of the telephone numbers that contain 919. You can search for any virtual member manager property because the command is generic.

Parameters and return values

-principalName

Specifies the principal name of the user that is used as the logon ID for the user in the system. This parameter maps to the principalName property in virtual member manager. You must specify only one of the following parameters: principalName, uid, cn, sn, or ibm-primaryEmail. (String, optional)

-uid

Specifies the unique ID value for the user for whom you want to search. This parameter maps to the uid property in virtual member manager. You must specify only one of the following parameters: principalName, uid, cn, sn, or ibm-primaryEmail. (String, optional)

-cn

Specifies the first name or given name of the user. This parameter maps to the cn property in virtual member manager. You must specify only one of the following parameters: principalName, uid, cn, sn, or ibm-primaryEmail. (String, optional)

-sn

Specifies the last name or family name of the user. This parameter maps to the sn property in virtual member manager. You must specify only one of the following parameters: principalName, uid, cn, sn, or ibm-primaryEmail. (String, optional)

-ibm-primaryEmail

Specifies the email address of the user. This parameter maps to the ibm-PrimaryEmail property in the virtual member manager. You must specify only one of the following parameters: principalName, uid, cn, sn, or ibm-primaryEmail. (String, optional)

-timeLimit

Specifies the maximum amount of time in milliseconds that the search can run. The default is not time limit. (String, optional)

-countLimit

Specifies the maximum number of results that you want returned from the search. By default, all users found in the search are returned. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask searchUsers {-principalName */IBM/US*}`
- Using Jython string:
`AdminTask.searchUsers (' [-principalName */IBM/US*]')`
- Using Jython list:
`AdminTask.searchUsers (['-principalName', '*/IBM/US*'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask searchUsers {-interactive}`
- Using Jython string:
`AdminTask.searchUsers (['-interactive'])`
- Using Jython list:
`AdminTask.searchUsers (['-interactive'])`

updateGroup

The **updateGroup** command updates the common name or the description of a group.

Parameters and return values

-uniqueName

Specifies the unique name value for the group for which you want to modify the properties. This parameter maps to the uniqueName property in virtual member manager. (String, required)

-cn

Specifies the new common name used for the group. This parameter maps to the cn property in virtual member manager. (String, optional)

-description

Specifies the new information about the group. This parameter maps to the description entity in a virtual member manager object. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask updateGroup {-uniqueName cn=opera tors,cn=groups,dc=yourco ,dc=com -cn groupA}`
- Using Jython string:
`AdminTask.updateGroup ('[-uniqueName cn=oper ators,cn=groups,dc=your co,dc=com -cn groupA]')`
- Using Jython list:
`AdminTask.updateGroup ([' -uniqueName', 'cn=oper ators,cn=groups,dc=yourco, dc=com', '-cn', 'groupA'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask updateGroup {-interactive}`
- Using Jython string:
`AdminTask.updateGroup ('[-interactive]')`
- Using Jython list:
`AdminTask.updateGroup (['-interactive'])`

updateUser

The **updateUser** command updates the following properties: uniqueName, uid, password, cn, sn, or ibm-primaryEmail.

Parameters and return values

-uniqueName

Specifies the unique name value for the user for which you want to modify the properties. This parameter maps to the uniqueName property in virtual member manager. (String, required)

-uid

Specifies the new unique ID value for the user. This parameter maps to the uid property in virtual member manager. (String, optional)

-password

Specifies the new password for the user. This parameter maps to the password property in virtual member manager. (String, optional)

-confirmPassword

Specifies the password again to validate how it was entered on the password parameter. This parameter maps to the password property in virtual member manager. (String, optional)

-cn

Specifies the new first name or given name of the user. This parameter maps to the cn property in virtual member manager. (String, optional)

-surname

Specifies the new last name or family name of the user. This parameter maps to the sn property in virtual member manager. (String, optional)

-ibm-primaryEmail

Specifies the new e-mail address of the user. This parameter maps to the mail property in virtual member manager. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask updateUser {-uniqueName uid=dmyers,cn=users,dc=yourco, dc=com -uid 123}`
- Using Jython string:
`AdminTask.updateUser ('[-uniqueName uid= dmyers,cn=users,dc= yourco,dc=com -uid 123]')`
- Using Jython list:
`AdminTask.updateUser (['-uniqueName', 'uid =dmyers,cn=users,dc= yourco,dc=com', '-uid', '123'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask updateUser {-interactive}`
- Using Jython string:
`AdminTask.updateUser ('[-interactive]')`
- Using Jython list:
`AdminTask.updateUser (['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

DescriptivePropCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the DescriptivePropCommands group can be used to create, delete, and manage key manager setting in your configuration.

The DescriptivePropCommands command group for the AdminTask object includes the following commands:

- “deleteDescriptiveProp”
- “getDescriptiveProp” on page 1231
- “listDescriptiveProp” on page 1231
- “modifyDescriptiveProp” on page 1232

deleteDescriptiveProp

The **deleteDescriptiveProp** command deletes key manager settings from the configuration.

Target object

None

Parameters and return values

-parentDataType
(String, required)

- parentClassName**
(String, required)
- parentScopeName**
(String, optional)
- name**
(String, required)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteDescriptiveProp {-interactive}`
- Using Jython string:
`AdminTask.deleteDescriptiveProp ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteDescriptiveProp (['-interactive'])`

getDescriptiveProp

The **getDescriptiveProp** command obtains information about key manager settings.

Target object

None

Parameters and return values

- parentDataType**
(String, required)
- parentClassName**
(String, required)
- parentScopeName**
(String, optional)
- name**
(String, required)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask getDescriptiveProp {-interactive}`
- Using Jython string:
`AdminTask.getDescriptiveProp ('[-interactive]')`
- Using Jython list:
`AdminTask.getDescriptiveProp (['-interactive'])`

listDescriptiveProp

The **listDescriptiveProp** command lists the key managers within a particular management scope.

Target object

None

Parameters and return values

-parentDataType
(String, required)

-parentClassName
(String, required)

-parentScopeName
(String, optional)

-displayObjectName
Set the value of this parameter to `true` to list the key manager objects within the scope. Set the value of this parameter to `false` to list the strings that contain the key manager name and management scope. (Boolean, optional)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask listDescrip tiveProp {-interactive}`
- Using Jython string:
`AdminTask.listDescrip tiveProp ('[-interact ive]')`
- Using Jython list:
`AdminTask.listDescrip tiveProp (['-interact ive'])`

modifyDescriptiveProp

The **modifyDescriptiveProp** command modifies the settings of an existing key manager.

Target object

None

Parameters and return values

-parentDataType
(String, required)

-parentClassName
(String, required)

-parentScopeName
(String, optional)

-name
(String, required)

-value
(String, optional)

-type
(String, optional)

-displayNameKey
(String, optional)

-nlsRangeKey
(String, optional)

- hoverHelpKey**
(String, optional)
- range**
(String, optional)
- inclusive**
(Boolean, optional)
- firstClass**
(Boolean, optional)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask modifyDescriptiveProp {-interactive}`
- Using Jython string:
`AdminTask.modifyDescriptiveProp ('[-interactive]')`
- Using Jython list:
`AdminTask.modifyDescriptiveProp (['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

ManagementScopeCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. Inbound and outbound management scopes represent opposing directions during the connection handshake process. The commands and parameters in the ManagementScopeCommands group can be used to create, delete, and list management scopes.

The ManagementScopeCommands command group for the AdminTask object includes the following commands:

- “deleteManagementScope”
- “getManagementScope” on page 1234
- “listManagementScopes” on page 1234

deleteManagementScope

The **deleteManagementScope** command deletes a management object from the configuration.

Target object

None

Parameters and return values

- scopeName

The name that uniquely identifies the management scope. (String, required)

Examples

Batch mode example usage:

Interactive mode example usage:

getManagementScope

The **getManagementScope** command displays the setting of a management scope object.

Target object

None

Parameters and return values

- scopeName

The name that uniquely identifies the management scope. (String, required)

Examples

Batch mode example usage:

Interactive mode example usage:

listManagementScopes

The **listManagementScopes** command lists the management scopes in the configuration.

Target object

None

Parameters and return values

- displayObjectName

Set the value to true to display the object names of the management scope. (Boolean, optional)

Examples

Batch mode example usage:

Interactive mode example usage:

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

AuthorizationGroupCommands command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the AuthorizationGroupCommands group can be used to create and manage authorization groups.

The AuthorizationGroupCommands command group for the AdminTask object includes the following commands:

- “addResourceToAuthorizationGroup”
- “createAuthorizationGroup” on page 1236
- “deleteAuthorizationGroup” on page 1237
- “listAuthorizationGroups” on page 1237
- “listAuthorizationGroupsForGroupID” on page 1238
- “listAuthorizationGroupsForUserID” on page 1239
- “listAuthorizationGroupsOfResource” on page 1239
- “listResourcesOfAuthorizationGroup” on page 1240
- “listResourcesForGroupID” on page 1241
- “listResourcesForUserID” on page 1241
- “mapGroupsToAdminRole” on page 1242
- “mapUsersToAdminRole” on page 1243
- “removeGroupsFromAdminRole” on page 1244
- “removeResourceFromAuthorizationGroup” on page 1245
- “removeUsersFromAdminRole” on page 1246

addResourceToAuthorizationGroup

The **addResourceToAuthorizationGroup** command adds a resource instance to an existing authorization group. A resource instance cannot belong to more than one authorization group.

Target object

None

Parameters and return values

- **authorizationGroupName**

The name of the authorization group. (String, required)

- **resourceName**

The name of the resource instance that you want to add to an authorization group. (String, required)

The resourceName parameter should be in the following format:

ResourceType=ResourceName

where `ResourceType` is one of the following values: `Application`, `Server`, `ServerCluster`, `Node`, `NodeGroup`

`ResourceName` is the name of the resource instance, for example, `server1`.

The following are example uses of the `resourceName` parameter:

- `Node=node1:Server=server1`

This example uniquely identifies `server1`. `node1` is required if another `server1` exists on a different node.

- `Application=app1`

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask addResourceToAuthorizationGroup {-authorizationGroupName groupName -resourceName Application=app1}
```
- Using Jython string:

```
AdminTask.addResourceToAuthorizationGroup('[-authorizationGroupName groupName -resourceName Application=app1]')
```
- Using Jython list:

```
AdminTask.addResourceToAuthorizationGroup(['-authorizationGroupName', 'groupName', '-resourceName', 'Application=app1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask addResourceToAuthorizationGroup {-interactive}
```
- Using Jython string:

```
AdminTask.addResourceToAuthorizationGroup ('[-interactive]')
```
- Using Jython list:

```
AdminTask.addResourceToAuthorizationGroup (['-interactive'])
```

createAuthorizationGroup

The **createAuthorizationGroup** command creates a new authorization group. When you create a new authorization group, no members are associated with it. Also, no user to administrative role mapping for the authorization table is associated with the authorization group.

Target object

None

Parameters and return values

- **authorizationGroupName**

The name of the authorization group that you want to create. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createAuthorizationGroup {-authorizationGroupName groupName}
```
- Using Jython string:

```
AdminTask.createAuthorizationGroup('[-authorizationGroupName groupName]')
```
- Using Jython list:

```
AdminTask.createAuthorizationGroup(['-authorizationGroupName', 'groupName'])
```

Interactive mode example usage:

- Using Jacl:
`$AdminTask createAuthorizationGroup -interactive`
- Using Jython string:
`AdminTask.createAuthorizationGroup ('[-interactive]')`
- Using Jython list:
`AdminTask.createAuthorizationGroup (['-interactive'])`

deleteAuthorizationGroup

The **deleteAuthorizationGroup** command deletes an existing authorization group. When you delete an authorization group, the authorization table that corresponds is also deleted.

Target object

None

Parameters and return values

- authorizationGroup Name

The name of the authorization group that you want to delete. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteAuthorizationGroup {-authorizationGroupName groupName}`
- Using Jython string:
`AdminTask.deleteAuthorizationGroup('[-authorizationGroupName groupName]')`
- Using Jython list:
`AdminTask.deleteAuthorizationGroup(['-authorizationGroupName', 'groupName'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask deleteAuthorizationGroup {-interactive}`
- Using Jython string:
`AdminTask.deleteAuthorizationGroup ('[-interactive]')`
- Using Jython list:
`AdminTask.deleteAuthorizationGroup (['-interactive'])`

listAuthorizationGroups

The **listAuthorizationGroups** command lists the existing authorization groups.

Target object

None

Parameters and return values

- Parameters: None
- Returns: A list of short names of all existing authorization groups. (String [])

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listAuthorizationGroups`
- Using Jython:
`AdminTask.listAuthorizationGroups()`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listAuthorizationGroups {-interactive}`
- Using Jython string:
`AdminTask.listAuthorizationGroups ('[-interactive]')`
- Using Jython list:
`AdminTask.listAuthorizationGroups (['-interactive'])`

listAuthorizationGroupsForGroupID

The **listAuthorizationGroupsForGroupID** command lists all of the authorization groups to which a given user group has access. This command lists the authorization groups and the granted roles for each authorization group. The group ID can be a short name or a fully qualified domain name if the LDAP user registry is being used. This command will list `cell` as a group if the user has cell level access.

Target object

None

Parameters and return values

- **groupid**
The ID of the user group. (String, required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listAuthorizationGroupsForGroupID {-groupid userGroupName}`
- Using Jython string:
`AdminTask.listAuthorizationGroupsForGroupID('[-groupid userGroupName]')`
- Using Jython list:
`AdminTask.listAuthorizationGroupsForGroupID(['-groupid', 'userGroupName'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listAuthorizationGroupsForGroupID {-interactive}`
- Using Jython string:
`AdminTask.listAuthorizationGroupsForGroupID ('[-interactive]')`
- Using Jython list:
`AdminTask.listAuthorizationGroupsForGroupID (['-interactive'])`

listAuthorizationGroupsForUserID

The **listAuthorizationGroupsForUserID** command lists all of the authorization groups to which a given user has access. This command lists the authorization groups and the granted roles for each authorization group. The user ID and the group ID can be a short name or a fully qualified domain name if the LDAP user registry is being used. This command will list `cell` as a group if the user has cell level access.

Target object

None

Parameters and return values

- **userid**

The ID of the user. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listAuthorizationGroupsForUserID{-userid userName}
```
- Using Jython string:

```
AdminTask.listAuthorizationGroupsForUserID(['-userid userName'])
```
- Using Jython list:

```
AdminTask.listAuthorizationGroupsForUserID(['-userid', 'userName'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listAuthorizationGroupsForUserID {-interactive}
```
- Using Jython string:

```
AdminTask.listAuthorizationGroupsForUserID (['-interactive'])
```
- Using Jython list:

```
AdminTask.listAuthorizationGroupsForUserID (['-interactive'])
```

listAuthorizationGroupsOfResource

The **listAuthorizationGroupsOfResource** command lists authorization groups for a given resource. If the value of the `traverseContainedObjects` parameter is `false`, only the authorization group of the resource is returned. If the value of the `traverseContainedObjects` parameter is `true`, it returns the authorization group of the resource and the authorization groups of all the parent resources in the containment tree.

Target object

None

Parameters and return values

- **resourceName**

The name of the resource. (String, required)

The `resourceName` parameter must be in the following format:

```
ResourceType=ResourceName
```

where `ResourceType` can be any one of the following values: `Application`, `Server`, `ServerCluster`, `Node`, or `NodeGroup`.

ResourceName is the name of the resource instance, for example, server1.

The following are examples of the resourceName parameter:

```
Node=node1:Server=server
```

This example uniquely identifies *server1*. The name of the node is required if a server on a different node uses the same server name.

```
Application=app1
```

- **traverseContained Resources**

Finds the authorization groups of all the parent resources by traversing the resource containment tree upwards. The default value is false. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listAuthorizationGroupsOfResource {-resourceName Application=app1}
```
- Using Jython string:

```
AdminTask.listAuthorizationGroupsOfResource(['-resourceName Application=app1'])
```
- Using Jython list:

```
AdminTask.listAuthorizationGroupsOfResource(['-resourceName', 'Application=app1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listAuthorizationGroupsOfResource {-interactive}
```
- Using Jython string:

```
AdminTask.listAuthorizationGroupsOfResource (['-interactive'])
```
- Using Jython list:

```
AdminTask.listAuthorizationGroupsOfResource (['-interactive'])
```

listResourcesOfAuthorizationGroup

The **listResourcesOfAuthorizationGroup** command lists all of the resources within the given authorization group.

Target object

None

Parameters and return values

- **authorizationGroupName**

The name of the authorization group. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listResourcesOfAuthorizationGroup {-authorizationGroupName groupName}
```
- Using Jython string:

```
AdminTask.listResourcesOfAuthorizationGroup(['-authorizationGroupName groupName'])
```
- Using Jython list:

```
AdminTask.listResourcesOfAuthorizationGroup(['-authorizationGroupName', 'groupName'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listResourcesOfAuthorizationGroup {-interactive}
```
- Using Jython string:

```
AdminTask.listResourcesOfAuthorizationGroup ('[-interactive]')
```
- Using Jython list:

```
AdminTask.listResourcesOfAuthorizationGroup (['-interactive'])
```

listResourcesForGroupID

The **listResourcesForGroupID** command lists all the objects that a given group has access to. This command lists the resources and the granted roles for each resource. The resources that this command returns include the resources from the authorization groups to which the user group is granted roles and the resources that are descendants of the resources with in authorization groups to which the user group is granted access to any role. The group ID can be a short name or fully qualified domain name if a LDAP user registry is used.

Target object

None

Parameters and return values

- **groupid**

The ID of the user group. (String, required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listResourcesForGroupID {-groupid userGroupName}
```
- Using Jython string:

```
AdminTask.listResourcesForGroupID ('[-groupid userGroupName']')
```
- Using Jython list:

```
AdminTask.listResourcesForGroupID (['-groupid', 'userGroupName'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listResourcesForGroupID {-interactive}
```
- Using Jython string:

```
AdminTask.listResourcesForGroupID ('[-interactive]')
```
- Using Jython list:

```
AdminTask.listResourcesForGroupID (['-interactive'])
```

listResourcesForUserID

The **listResourcesForUserID** command lists all the objects that a given user has access to. This command lists the resources and the granted roles for each resource. The resources that this command returns include the resources from the authorization groups to which the user is granted roles and the

resources that are descendants of the resources with in authorization groups to which the user is granted access to any role. The user ID can be a short name or fully qualified domain name if a LDAP user registry is used.

Target object

None

Parameters and return values

- userid

The ID of the user. (String, required).

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask listResourcesForUserID {-userid userName }`
- Using Jython string:
`AdminTask.listResourcesForUserID(['-userid userName'])`
- Using Jython list:
`AdminTask.listResourcesForUserID(['-userid', 'userName'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask listResourcesForUserID {-interactive}`
- Using Jython string:
`AdminTask.listResourcesForUserID ('[-interactive]')`
- Using Jython list:
`AdminTask.listResourcesForUserID (['-interactive'])`

Example output:

```
{deployer=[], operator=[], administrator=[cells/IBM-LP1 6L31HVE8Cell07/clusters/C1| cluster.xml,
cells/IBM-LP16L 31HVE8Cell07/nodes/IBM-LP16L 31HVE8Node05/servers/cm1|ser ver.xml],
monitor=[], configurator=[]}
```

mapGroupsToAdminRole

The **mapGroupsToAdminRole** command maps group IDs to one or more administrative roles in an authorization group. The name of the authorization group that you provide determines which authorization table will be used. If you do not specify an authorization group name, the mapping is done to the cell level authorization table. The group ID can be a short name or a fully qualified domain name if the LDAP user registry is used.

Target object

None

Parameters and return values

- authorizationGroup Name

The name of the authorization group. If you do not specify this parameters, the cell level authorization group is assumed. (String, optional)

- roleName

The name of the administrative role. (String, required)

- groupids

The list of group IDs that will mapped to the administrative role. (String[], required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask mapGroupsToAdminRole {-authorizationGroupName groupName - roleName administrator -groupids group1}`
- Using Jython string:
`AdminTask.mapGroupsToAdminRole(['-authorizationGroupName groupName -roleName administrator -groupids group1'])`
- Using Jython list:
`AdminTask.mapGroupsToAdminRole(['-authorizationGroupName', 'groupName', '-roleName', 'administrator', '-groupids', 'group1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask mapGroupsToAdminRole {-interactive}`
- Using Jython string:
`AdminTask.mapGroupsToAdminRole (['-interactive'])`
- Using Jython list:
`AdminTask.mapGroupsToAdminRole (['-interactive'])`

mapUsersToAdminRole

The **mapUsersToAdminRole** command maps user IDs to one or more administrative roles in the authorization group. The name of the authorization group that you provide determines the authorization table. If you do not specify the name of the authorization group, the mapping is done to the cell level authorization table. The user ID can be a short name or fully qualified domain name in case LDAP user registry is used.

Target object

None

Parameters and return values

- authorizationGroup Name

The name of the authorization group. If you do not specify this parameter, the cell level authorization group is assumed. (String, optional)

- roleName

The name of the administrative role. (String, required)

- userids

The list of user IDs that will be mapped to the administrative role (String[], required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask mapUsersToAdminRole {-authorizationGroupName groupName - roleName administrator -userids user1}`
- Using Jython string:

```
AdminTask.mapUsersToAdminRole(['-authorizationGroupName groupName -roleName administrator -userids user1'])
```

- Using Jython list:

```
AdminTask.mapUsersToAdminRole(['-authorizationGroupName', 'groupName', '-roleName', 'administrator',  
'-userids', 'user1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask mapUsersToAdminRole {-interactive}
```

- Using Jython string:

```
AdminTask.mapUsersToAdminRole ('[-interactive]')
```

- Using Jython list:

```
AdminTask.mapUsersToAdminRole (['-interactive'])
```

removeGroupsFromAdminRole

The **removeGroupsFromAdminRole** command removes previously mapped group IDs from administrative roles in the authorization group. The name of the authorization group that you provide determines which authorization table is involved. If you do not specify an authorization group name, the group IDs are removed from the cell level authorization table. The group ID can be a short name or fully qualified domain name if a LDAP user registry is used.

Target object

None

Parameters and return values

- authorizationGroupName

The name of the authorization group. If you do not specify this parameter, the cell level authorization group is assumed. (String, optional)

- roleName

The name of the administrative role. (String, required)

- userids

A list of group IDs that you want to remove from the administrative role. (String[], required)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask removeGroupsFromAdminRole {-authorizationGroupName groupName - roleName administrator -groupids group1}
```

- Using Jython string:

```
AdminTask.removeGroupsFromAdminRole(['-authorizationGroupName groupName -roleName administrator -groupids group1'])
```

- Using Jython list:

```
AdminTask.removeGroupsFromAdminRole(['-authorizationGroupName', 'groupName', '-roleName', 'administrator',  
'-groupids', 'group1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask removeGroupsFromAdminRole {-interactive}
```

- Using Jython string:

```
AdminTask.removeGroupsFromAdminRole ('[-interactive]')
```

- Using Jython list:

```
AdminTask.removeGroupsFromAdminRole (['-interactive'])
```

removeResourceFromAuthorizationGroup

The **removeResourceFromAuthorizationGroup** command removes resources from an existing authorization group. If you do not specify the authorization group, it will be determined and the resource will be removed from that authorization group.

Target object

None

Parameters and return values

- authorizationGroup Name

The name of the authorization group. (String, optional)

- resourceName

The name of the resource instance that you want to remove from the authorization group. (String, required)

The resourceName parameter must be in the following format:

```
ResourceType=ResourceName
```

where the ResourceType can be any of the following: Application, Server, ServerCluster, Node, or NodeGroup.

The ResourceName is the name of the resource instance, for example, server1.

The following are examples of the resourceName parameter:

```
Node=node1:Server=server1
```

This example uniquely identifies server1. node1 is required if the name of the server exists on multiple nodes.

```
Application=app1
```

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask removeResourceFromAuthorizationGroup {-authorizationGroupName groupName -resourceName Application=app1}
```

- Using Jython string:

```
AdminTask.removeResourceFromAuthorizationGroup(['-authorizationGroupName groupName -resourceName Application=app1'])
```

- Using Jython list:

```
AdminTask.removeResourceFromAuthorizationGroup(['-authorizationGroupName', 'groupName', '-resourceName', 'Application=app1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask removeResourceFromAuthorizationGroup {-interactive}
```

- Using Jython string:

```
AdminTask.removeResourceFromAuthorizationGroup (['-interactive'])
```

- Using Jython list:

```
AdminTask.removeResourceFromAuthorizationGroup (['-interactive'])
```

removeUsersFromAdminRole

The **removeUsersFromAdminRole** command removes previously mapped user IDs from administrative roles in the authorization group. The name of the authorization group that you provide determines which authorization table is involved. If you do not specify an authorization group name, the user ID from the cell level authorization table will be used. The user ID can be a short name or a fully qualified domain name if a LDAP user registry is used.

Target object

None

Parameters and return values

- authorizationGroup Name

The name of the authorization group. If you do not specify this parameter, the cell level authorization group is assumed. (String, optional)

- roleName

The name of the administrative role. (String, required)

- userids

A list of user IDs that you want to remove from the administrative role. (String[], required)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask removeUsersFromAdminRole {-authorizationGroupName groupName - roleName administrator -userids user1}`
- Using Jython string:
`AdminTask.removeUsersFromAdminRole(['-authorizationGroupName groupName -roleName administrator -userids user1'])`
- Using Jython list:
`AdminTask.removeUsersFromAdminRole(['-authorizationGroupName', 'groupName', '-roleName', 'administrator', '-userids', 'user1'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask removeUsersFromAdminRole {-interactive}`
- Using Jython string:
`AdminTask.removeUsersFromAdminRole (['-interactive'])`
- Using Jython list:
`AdminTask.removeUsersFromAdminRole (['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

ChannelFrameworkManagement command group for the AdminTask object

You can use the Jython or Jacl scripting languages to configure security. The commands and parameters in the ChannelFrameworkManagement group can be used to create and manage transport channels and transport channel chains.

The ChannelFrameworkManagement command group for the AdminTask object includes the following commands:

- “createChain”
- “deleteChain” on page 1248
- “listChainTemplates” on page 1249
- “listChains” on page 1249

createChain

The **createChain** command creates a new chain of transport channels that are based on a chain template.

Target object

The instance of the transport channel service under which the new chain is created. (ObjectName, required)

Required parameters and return values

- template

The chain template on which to base the new chain. (ObjectName, required)

- name

The name of the new chain. (String, required)

- endPoint

The name of the end point to be used by the instance of the TCP inbound channel in the new chain if the chain is an inbound chain. (ObjectName, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask createChain (cells/ rohitbuildCell101/nodes/rohit buildCellManager01/servers/
dmgr|server.xml#TransportChannelService_1) {-template WebContainer(templates/ chains|webcontainer-chains.xml #Chain_1)
-name trialChain1}
```

```
$AdminTask createChain (cells/ rohitbuildCell101/nodes/rohit buildCellManager01/servers/dmgr
|server.xml#TransportChannelService_1) {-template WebContainer(templates/ chains|webcontainer-chains.xml# Chain_1)
-name trialChain1 -endPoint (cells/rohitbuild Cell101/nodes/rohitbuildCellMa nager01|serverindex.xml#End Point_3) }
```

- Using Jython string:

```
AdminTask.createChain('cells/ rohitbuildCell101/nodes/rohitbu ildCellManager01/servers/dmgr|
server.xml#TransportChannelSer vice_1', '[-template "WebConta iner(templates/chains|webconta iner-chains.xml#Chain_1)" -name
trialChain]')
```

```
AdminTask.createChain('cells/rohitbuildCell01/nodes/rohitbuildCellManager01/servers/dmgr
|server.xml#TransportChannelService_1', ['-template "WebContainer(templates/chains|webcontainer-chains.xml#Chain_1)" -name
trialChain -endPoint "(cells/rohitbuildCell01/nodes/rohitbuildCellManager01|serverindex.xml#EndPoint_3)"]')
```

- Using Jython list:

```
AdminTask.createChain('cells/rohitbuildCell01/nodes/rohitbuildCellManager01/serve
rs/dmgr|server.xml#TransportChannelService_1', ['-template', "WebContainer(templates/chains|webcontaine
r-chains.xml#Chain_1)", '-name', 'trialChain'])
```

```
AdminTask.createChain('cells/rohitbuildCell01/nodes/rohitbuildCellManager01/servers/
dmgr|server.xml#TransportChannelService_1', ['-template', "WebContainer(templates/chains|webcontainer-chains.xml#Cha
in_1)",
'-name', 'trialChain', '-endPoint', "(cells/rohitbuildCell01/nodes/rohitbuildCellManager01|serverindex.xml#
EndPoint_3)"])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask createChain {-interactive}
```

- Using Jython string:

```
AdminTask.createChain ('[-interactive]')
```

- Using Jython list:

```
AdminTask.createChain (['-interactive'])
```

deleteChain

The **deleteChain** command deletes an existing chain and, optionally, the transport channels in the chain.

Target object

The chain to be deleted. (Object name, required)

Required parameters and return values

- deleteChannels

If the value of this attribute is true, non-shared transport channels used by the specified chain will be deleted. (Boolean, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask deleteChain trialChain1 (cells/rohitbuildCell01/nodes/roh
itbuildCellManager01/servers/dmgr | server.xml#Chain_1093554462922)
$AdminTask deleteChain trialChain (cells/rohitbuildCell01/nodes/roh
itbuildCellManager01/servers/dmgr |server.xml#Chain_1093554378078) {-deleteChannels true}
```

- Using Jython string:

```
AdminTask.deleteChain('trialChain1(cells/rohitbuildCell01/nodes/roh
itbuildCellManager01/servers/dmgr| server.xml#TransportChannelService_1)')
AdminTask.deleteChain('trialChain1 (cells/rohitbuildCell01/nodes/rohi
tbuildCellManager01/servers/dmgr| server.xml#TransportChannelService_1)', ['-deleteChannels true'])
```

- Using Jython list:

```
AdminTask.deleteChain('trialChain1 (cells/rohitbuildCell01/nodes/roh
itbuildCellManager01/servers/dmgr| server.xml#TransportChannelService_1)')
AdminTask.deleteChain('trialChain1 (cells/rohitbuildCell01/nodes/rohi
tbuildCellManager01/servers/dmgr| server.xml#TransportChannelService_1)', ['-deleteChannels', 'true'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteChain {-interactive}
```

- Using Jython string:

```
AdminTask.deleteChain ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteChain (['-interactive'])
```

listChainTemplates

The **listChain Templates** command displays a list of templates that you can use to create chains in this configuration. All templates have a certain type of transport channel as the last transport channel in the chain.

Target object

None

Required parameters and return values

- acceptorFilter

The templates returned by this method all have a transport channel instance of the specified type as the last transport channel in the chain. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listChainTemplates {}  
$AdminTask listChainTemplates "-acceptorFilter WebContainer InboundChannel"
```

- Using Jython string:

```
AdminTask.listChainTemplates()  
AdminTask.listChainTemplates (['-acceptorFilter WebCont ainerInboundChannel'])
```

- Using Jython list:

```
AdminTask.listChainTemplates()  
AdminTask.listChainTemplates (['-acceptorFilter', 'WebC ontainerInboundChannel'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listChainTemplates {-interactive}
```

- Using Jython string:

```
AdminTask.listChainTemplates (['-interactive'])
```

- Using Jython list:

```
AdminTask.listChainTemplates (['-interactive'])
```

listChains

The **listChains** command lists all the chains that are configured under a particular instance of the transport channel service.

Target object

The instance of the transport channel service under which the chains are configured. (ObjectName, required)

Required parameters and return values

- acceptorFilter

The chains that are returned by this parameter will have a transport channel instance of the type that you specify as the last transport channel in the chain. (String, optional)

- endPointFilter:

The chains returned by this parameter will have a TCP inbound channel using an end point with the name that you specify.(String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask listChains (cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328)
$AdminTask listChains (cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328) {-acceptorFilter WebContai nerInboundChannel}
$AdminTask listChains (cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328) {-end PointFilter WC_adminhost}
```

- Using Jython string:

```
AdminTask.listChains('(cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328)')
AdminTask.listChains('(cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328)', ['-acceptorFilter WebContai nerInboundChannel'])
AdminTask.listChains('(cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328)', ['-endPointFilter WC_adminhost'])
```

- Using Jython list:

```
AdminTask.listChains('(cells/ rohitbuildCell101/nodes/rohit buildNode01/servers/server2|
server.xml#TransportChannel Service_1093445762328)')
AdminTask.listChains('(cells /rohitbuildCell101/nodes/rohit itbuildNode01/servers/server 2
|server.xml#TransportChanne lService_1093445762328)', ['-acceptorFilter', 'WebCon tainerInboundChannel'])
AdminTask.listChains('(cells /rohitbuildCell101/nodes/roh itbuildNode01/servers/server
2|server.xml #TransportChanne lService_1093445762328)', ['-endPointFilter', 'WC_admi nhost'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask listChains {-interactive}
```

- Using Jython string:

```
AdminTask.listChains ('[-interactive]')
```

- Using Jython list:

```
AdminTask.listChains (['-interactive'])
```

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

SpnegoTAICommands group for the AdminTask object (deprecated)

You can use the Jython or Jacl scripting languages to configure security with the wsadmin tool. The commands and parameters in the SpnegoTAICommands group can be used to create and manage configurations that are used by the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI).

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application

Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

The SpnegoTAICommands command group for the AdminTask object includes the following commands:

- “addSpnegoTAIProperties”
- “deleteSpnegoTAIProperties” on page 1252
- “modifySpnegoTAIProperties” on page 1253
- “showSpnegoTAIProperties” on page 1254
- “createKrbConfigFile” on page 1254

addSpnegoTAIProperties

The **addSpnego TAI Properties** command adds properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for the application server.

Target object

None

Parameters and return values

-spnId

This is the SPN identifier for the group of custom properties that are to be defined with this command. If you do not specify this parameter, an unused SPN identifier is assigned. (String, optional)

-host

Specifies the host name portion in the SPN used by the SPNEGO TAI to establish a Kerberos secure context. (String, required)

-filter

Defines the filtering criteria used by the class specified with the above attribute. If no filter is specified, all HTTP requests are subject to SPNEGO authentication. (String, optional)

-filterClass

Specifies the name of the Java class used by the SPNEGO TAI to select which HTTP requests will be subject to SPNEGO authentication. If no filter class is specified, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used. (String, optional)

-noSpnegoPage

Specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application if it does not support SPNEGO authentication. (String, optional).

If you do not specify the noSpnegoPage attribute then the default is used:

```
"<html><head><title> SPNEGO authentication is not supported. </title></head>" +  
"<body>SPNEGO authentication is not supported on this client.</body> </html>";
```

-ntlmTokenPage

Specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application when the SPNEGO token received by the interceptor after the challenge-response handshake contains a NT LAN manager (NTLM) token instead of the expected SPNEGO token. (String, optional).

If you do not specify the ntlmTokenPage attribute then the default is used:

```
"<html><head><title> An NTLM Token was received.</title> </head>" + "<body>Your browser configuration  
is correct, but you have not logged into a supported Windows Domain." + "<p>Please login to the application  
using the normal login page.</html>";
```

-trimUserName

Specifies whether (`true`) or not (`false`) the SPNEGO TAI is to remove the suffix of the principal user name, starting from the @ that precedes the Kerberos realm name. If this attribute is set to `true`, the suffix of the principal user name is removed. If this attribute is set to `false`, the suffix of the principal name is retained. The default value used is `true`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask addSpnegoTAIProperties -host myhost.ibm.com -filter user-agent%=IE 6`
- Using Jython string:
`AdminTask.addSpnegoTAIProperties ('[-host myhost.ibm.com -filter user-agent%=IE 6]')`
- Using Jython list:
`AdminTask.addSpnegoTAIProperties (['-host', 'myhost.ibm.com', '-filter', 'user-agent%=IE', '6'])`

Interactive mode example usage:

- Using Jacl:
`$AdminTask addSpnegoTAIProperties -interactive`
- Using Jython string:
`AdminTask.addSpnegoTAIProperties (['-interactive'])`
- Using Jython list:
`AdminTask.addSpnegoTAIProperties ['-interactive'])`

deleteSpnegoTAIProperties

The **deleteSpnegoTAIProperties** command deletes properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Target object

None

Parameters and return values

-spnId

The SPN identifier for the group of custom properties that are to be deleted with this command. If you do not specify this parameter, all SPNEGO TAI custom properties are deleted. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:
`$AdminTask deleteSpnegoTAIProperties {-spnId 2}`
- Using Jython string:
`AdminTask.deleteSpnegoTAIProperties (['-spnId 2'])`
- Using Jython list:
`AdminTask.deleteSpnegoTAIProperties (['-spnId', '2'])`

Interactive mode example usage:

- Using Jacl:

```
$AdminTask deleteSpnegoTAI Properties -interactive
```

- Using Jython string:

```
AdminTask.deleteSpnegoTAI Properties ('[-interactive]')
```

- Using Jython list:

```
AdminTask.deleteSpnegoTAI Properties ['-interactive'])
```

modifySpnegoTAIProperties

The **modifySpnego TAIProperties** command modifies the properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Target object

None

Parameters and return values

-spnId

The SPN identifier for the group of custom properties that are to be defined with this command. (String, required)

-host

Specifies the host name portion in the SPN used by the SPNEGO TAI to establish a Kerberos secure context. (String, optional)

-filter

Defines the filtering criteria used by the class specified with the above attribute. (String, optional)

-filterClass

Specifies the name of the Java class used by the SPNEGO TAI to select which HTTP requests will be subject to SPNEGO authentication. If no class is specified, all HTTP requests will be subject to SPNEGO authentication. (String, optional)

-noSpnegoPage

Specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application if it does not support SPNEGO authentication. (String, optional)

-ntlmTokenPage

Specifies the URL of a resource that contains the content the SPNEGO TAI will include in the HTTP response to be displayed by the (browser) client application when the SPNEGO token received by the interceptor after the challenge-response handshake contains a NT LAN manager (NTLM) token instead of the expected SPNEGO token. (String, optional)

-trimUserName

Specifies whether (`true`) or not (`false`) the SPNEGO TAI is to remove the suffix of the principal user name, starting from the "@" that precedes the Kerberos realm name. If this attribute is set to `true`, the suffix of the principal user name is removed. If this attribute is set to `false`, the suffix of the principal name is retained. The default value used is `true`. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask modifySpnegoTAI PROPERTIES -spnId 1 -filter host==myhost.company.com
```

- Using Jython string:

```
AdminTask.modifySpnegoTAI PROPERTIES ('[-spnId 1 -filter host==myhost.com pany.com]')
```

- Using Jython list:

```
AdminTask.modifySpnegoTAI PROPERTIES (['-spnId', '1', '-filter', 'host=my host.company.com'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask modifySpnegoTAI Properties -interactive
```

- Using Jython string:

```
AdminTask.modifySpnegoTAI Properties ('[-interactive]')
```

- Using Jython list:

```
AdminTask.modifySpnegoTAI Properties ['-interactive'])
```

showSpnegoTAIProperties

The **showSpnego TAI Properties** command displays the properties in the configuration of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Target object

None

Parameters and return values

-spnId

The service principal name (SPN) identifier for the group of custom properties that are to be displayed with this command. If you do not specify this parameter, all SPNEGO TAI custom properties are displayed. (String, optional)

Examples

Batch mode example usage:

- Using Jacl:

```
$AdminTask showSpnegoTAI Properties -spnId 1
```

- Using Jython string:

```
AdminTask.showSpnegoTAI Properties ('[-spnId 1]')
```

- Using Jython list:

```
AdminTask.showSpnegoTAI Properties (['-spnId', '1'])
```

Interactive mode example usage:

- Using Jacl:

```
$AdminTask showSpnegoTAI Properties -interactive
```

- Using Jython string:

```
AdminTask.showSpnegoTAI Properties (['-interact ive'])
```

- Using Jython list:

```
AdminTask.showSpnegoTAI Properties ['-interact ive'])
```

createKrbConfigFile

The **createKrb ConfigFile** command creates the Kerberos configuration file for use with the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Target object

None

Parameters and return values

-krbPath

Provides the fully qualified file system location of the Kerberos configuration (krb5.ini or krb5.conf) file. (String, required)

-realm

Provides the Kerberos realm name. The value of this attribute is used by the SPNEGO TAI to form the Kerberos service principal name for each of the hosts specified with the property com.ibm.ws.security.spnego.SPN<id>.hostname (String, required)

-kdcHost

Provides the host name of the Kerberos Key Distribution Center (KDC). (String, required)

-kdcPort

Provides the port number of the KDC. The default value, if not specified, is 88. (String, optional)

-dns

Provides the default domain name service (DNS) that is used to produce a fully qualified host name. (String, required)

-keytabPath

Provides the file system location of the Kerberos keytab file. (String, required)

-encryption

Identifies the list of supported encryption types, separated by a space. The specified value is used for the default_tkt_encypes and default_tgs_encypes. The default encryption types, if not specified, are des-cbc-md5 and rc4-hmac. (String, optional)

Examples

Interactive mode example usage:

- Using Jacl:
`$AdminTask createKrbCo nfigFile -interactive`
- Using Jython string:
`AdminTask.createKrbCon figFile ['-interactive']`
- Using Jython list:
`AdminTask.createKrbCon figFile ['-interactive'])`

Related tasks

Using the AdminTask object for scripted administration

Use the AdminTask object to access a set of administrative commands that provide an alternative way to access the configuration commands and the running object management commands.

Related reference

Commands for the AdminTask object

Use the AdminTask object to run administrative commands with the wsadmin tool.

The Kerberos configuration file

The Kerberos configuration properties, krb5.ini or krb5.conf files, must be configured on every WebSphere Application Server instance in a cell in order to use the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) for WebSphere Application Server.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

The default Kerberos configuration file name for Windows is krb5.ini,. For other platforms is the default Kerberos configuration file name is krb5.conf,. The default location for the Kerberos configuration file is shown below:

Operating System	Default Location
Windows	c:\winnt\krb5.ini Note: If the krb5.ini file is not located in the c:\winnt directory it might be located in c:\windows directory.
Linux	/etc/krb5.conf
other UNIX-based	/etc/krb5/krb5.conf
z/OS	/etc/krb5/krb5.conf
i5/OS	/QIBM/UserData/OS400/NetworkAuthentication/krb5.conf

Note: If you don't use the default location and Kerberos configuration file name, then you have to update *.krb5ConfigFile properties in the soap.client.props, ipc.client.props and sas.client.props files. Also, if the client programmatic login uses the WSKRBLLogin module, you must also set the java.security.krb5.conf JVM property.

For SPNEGO TAI, if you don't use the default location and Kerberos configuration file name, then you must specify the java.security.krb5.conf JVM property.

The default Kerberos configuration file on Windows is /winnt/krb5.ini and on a distributed environment is /etc/krb5. If you specify another location path, then you must also specify the java.security.krb5.conf JVM property.

For example, if your krb5.conf file is specified at /opt/IBM/WebSphere/profiles/AppServer/etc/krb5.conf, then you need to specify -Djava.security.krb5.conf=/opt/IBM/WebSphere/profiles/AppServer/etc/krb5.conf.

The WebSphere runtime code searches for the Kerberos configuration file in the order as follows:

1. The file referenced by the Java property java.security.krb5.conf
2. <java.home>/lib/security/krb5.conf
3. c:\winnt\krb5.ini on Microsoft Windows platforms
4. /etc/krb5/krb5.conf on UNIX platforms
5. /etc/krb5.conf on Linux™ platforms.

Use the wsadmin utility to configure the SPNEGO TAI for WebSphere Application Server:

1. Start WebSphere Application Server.
2. Start the command-line utility by running the **wsadmin** command from the *app_server_root/bin* directory.
3. At the **wsadmin** prompt, enter the following command:

```
$AdminTask createKrbConfigFile
```

You can use the following parameters with this command:

Option	Description
<krbPath>	This parameter is required. It provides the fully qualified file system location of the Kerberos configuration (krb5.ini or krb5.conf) file.
<realm>	This parameter is required. It provides the Kerberos realm name. The value of this attribute is used by the SPNEGO TAI to form the Kerberos service principal name for each of the hosts specified with the property com.ibm.ws.security.spnego.SPN<id>.hostName.
<kdcHost>	This parameter is required. It provides the host name of the Kerberos Key Distribution Center (KDC).
<kdcPort>	This parameter is optional. It provides the port number of the KDC. The default value, if not specified, is 88.
<dns>	This parameter is required. It provides the default domain name service (DNS) that is used to produce a fully qualified host name.
<keytabPath>	This parameter is required. It provides the file system location of the Kerberos keytab file.
<encryption>	This parameter is optional. It identifies the list of supported encryption types, separated by a space. The specified value is used for the default_tkt_encetypes and default_tgs_encetypes. The default encryption types, if not specified, are des-cbc-md5 and rc4-hmac.

In the following example, the wsadmin command creates the krb5.ini file in the c:\winnt directory. The default Kerberos keytab file is also in c:\winnt. The actual Kerberos realm name is WSSEC.AUSTIN.IBM.COM and the KDC host name is host1.austin.ibm.com.

```
wsadmin>$AdminTask createKrbConfigFile {-krbPath
c:\winnt\krb5.ini -realm WSSEC.AUSTIN.IBM.COM -kdcHost host1.austin.ibm.com
-dns austin.ibm.com -keytabPath c:\winnt\krb5.keytab}
```

The wsadmin command above creates a krb5.ini file as follows:

```
[libdefaults]
default_realm = WSSEC.AUSTIN.IBM.COM
default_keytab_name = FILE:c:\winnt\krb5.keytab
default_tkt_encetypes = des-cbc-md5 rc4-hmac
default_tgs_encetypes = des-cbc-md5 rc4-hmac
[realms]
WSSEC.AUSTIN.IBM.COM = {
kdc = host1.austin.ibm.com:88
default_domain = austin.ibm.com
}
[domain_realm]
.austin.ibm.com = WSSEC.AUSTIN.IBM.COM
```

Note:

- A Kerberos keytab file contains a list of keys that are analogous to user passwords. It is important for hosts to protect their Kerberos keytab files by storing them on the local disk. The krb5.conf file permission must be 644, which means that you can read and write the file; however, members of the group that the file belongs to, and all others can only read the file.
- If the run time cannot read the default_tkt_encetypes or default_tgs_encetypes entries in the krb5.ini file, their values are missing, or their values are not supported, the DES-CBC-MD5 value is used by default.

Kerberos configuration settings, the Kerberos key distribution center (KDC) name, and realm settings for the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI) are provided in the Kerberos configuration file or through `java.security.krb5.kdc` and `java.security.krb5.realm` system property files.

SPNEGO Web authentication configuration commands

Use `wsadmin` commands to configure, unconfigure, validate, or display Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) in the security configuration.

Configure SPNEGO Web authentication

Note: You must first have a workable Kerberos configuration file and a Kerberos keytab file. Read about “Creating a Kerberos configuration file” on page 270 and “Creating a Kerberos service principal and keytab file” on page 273 for more information.

Use the `configureSpnego` command to configure SPNEGO as a Web authenticator in the security configuration.

At the `wsadmin` prompt, enter the following command for help:

```
Wsadmin>$AdminTask help configureSpnego
```

You can use the following parameters with the `configureSpnego` command:

Option	Description
<enabled>	This parameter is optional. It enables SPNEGO Web authentication.
<dynamicReload>	This parameter is optional. It enables dynamic reload of SPNEGO Web authentication filters.
<allowAppAuthMethodFallback>	This parameter is optional. It allows fall back to the application authentication mechanism.
<krb5Config>	This parameter is required. It supplies the directory location and file name of the configuration (<code>krb5.ini</code> or <code>krb5.conf</code>) file.
<krb5Keytab>	This parameter is optional. It supplies the directory location and file name of the Kerberos keytab file. If you do not specify this parameter, the default keytab in the Kerberos configuration file is used.

Unconfigure SPNEGO Web authentication

Use the `unconfigureSpnego` command to unconfigure SPNEGO Web authentication in the security configuration.

At the `wsadmin` prompt, enter the following command for help:

```
wsadmin>$AdminTask help unconfigureSpnego
```

Show SPNEGO Web authentication

Use the `showSPNEGO` command to display the SPNEGO Web authentication in the security configuration.

At the `wsadmin` prompt, enter the following command for help:

```
wsadmin>$AdminTask help showSpnego
```

Validate Kerberos configuration

Use the `validateKrbConfig` command to validate the Kerberos configuration data either in the global security file `security.xml` or specified as an input parameter.

At the **wsadmin** prompt, enter the following command for help:

```
wsadmin>$AdminTask help validateKrbConfig
```

You can use the following parameters with the `validateKrbConfig` command:

Option	Description
<checkConfigOnly>	Checks the Kerberos configuration without validating, You must use global security for this check.
<useGlobalSecurityConfig>	Uses the Global Security configuration data, <code>security.xml</code> , instead of input parameters.
<validateKrbRealm>	Validates the Kerberos realm against the default Kerberos realm in the Kerberos configuration file (<code>krb5.ini</code> or <code>krb5.conf</code>).
<serverId>	Specifies the server identity that is used for internal process communications.
<serverIdPassword>	Specifies the password that is used for the server identity.
<krb5Spn>	Specifies the Kerberos service principal name in the Kerberos keytab file.
<krb5Config >	This parameter is required. It supplies the directory location and file name of the configuration (<code>krb5.ini</code> or <code>krb5.conf</code>) file.
<krb5Keytab>	This parameter is optional. It supplies the directory location and file name of the Kerberos keytab file. If you do not specify this parameter, the default keytab in the Kerberos configuration file is used.
<krb5Realm >	This parameter is required. It specifies the value for the Kerberos realm name.

SPNEGO Web authentication filter commands

Use **wsadmin** commands to add, modify, delete, or show Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) Web authentication filters in the security configuration.

Add SPNEGO web authentication filter

Use the **addSpnegoFilter** command to add a new SPNEGO Web authentication filter in the security configuration.

At the **wsadmin** prompt, enter the following command for help:

```
wsadmin>$AdminTask help addSpnegoFilter
```

You can use the following parameters with the **addSpnegoFilter** command:

Option	Description
<hostName>	This parameter is required. Use to supply a fully-qualified host name.
<krb5Realm>	This parameter is not required. Use to supply a Kerberos realm name. If the krb5Realm parameter is not specified, the default Kerberos realm name in the Kerberos configuration file is used.
<filterCriteria>	This parameter is not required. Use to supply the HTTP request filter rules. If the filterCriteria parameter is not specified, all of the HTTP requests are authenticated by SPNEGO.
<filterClass>	This parameter is not required. Use to supply the HTTP request filter rules. If the filterClass parameter is not specified, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used.
<trimUserName>	This parameter is not required. Use to indicate whether the Kerberos realm name is to be removed from the Kerberos principal name.
<enabledGssCredDelegate>	This parameter is not required. Use to indicate whether to extract and place the client GSS delegation credential in the subject. The default value is true.
<spnegoNotSupportedPage>	This parameter is not required. Use to supply the uniform resource identifier (URI) of the resource with a response to be used when SPNEGO is not supported. If this parameter is not specified, the default SPNEGO not supported error page is used.
<ntlmTokenReceivedPage>	This parameter is not required. Use to supply the URI of the resource with a response to be used when an NT LAN manager (NTLM) token is received. If this parameter is not specified, the default NTLM token received error page is used.

The following is an example of the **addSpnegoFilter** command:

```
wsadmin>$AdminTask addSpnegoFilter {
  -hostName ks.austin.ibm.com
  -krb5Realm WSSEC.AUSTIN.IBM.COM}
```

Modify SPNEGO web authentication filter

Use the **modifySpnegoFilter** command to modify SPNEGO filter attributes in the security configuration.

At the **wsadmin** prompt, enter the following command for help:

```
wsadmin>$AdminTask help modifySpnegoFilter
```

You can use the following parameters with the **modifySpnegoFilter** command:

Option	Description
<hostName>	This parameter is required. Use to supply a long host name. The hostname is an identifier, so you can not modify the hostname.

Option	Description
<krb5Realm>	This parameter is not required. Use to supply a Kerberos realm name. If the krb5Realm parameter is not specified, the default Kerberos realm name in the Kerberos configuration file is used.
<filterCriteria>	This parameter is not required. Use to supply the HTTP request filter rules. If the filterCriteria parameter is not specified, all of the HTTP requests are authenticated by SPNEGO. Note: Read about “Enabling and configuring SPNEGO Web authentication using the administrative console” on page 307 for more information about filter criteria.
<filterClass>	This parameter is not required. Use to supply the HTTP request filter rules. If the filterClass is not specified, the default filter class, com.ibm.ws.security.spnego.HTTPHeaderFilter, is used.
<trimUserName>	This parameter is not required. Use to indicate whether the Kerberos realm name is to be removed from the Kerberos principal name.
<enabledGssCredDelegate>	This parameter is not required. Use to indicate whether to extract and place the client GSS delegation credential in the subject. The default value is true.
<spnegoNotSupportedPage>	This parameter is not required. Use to supply the URI of the resource with a response to be used when SPNEGO is not supported. If this parameter is not specified, the default SPNEGO not supported error page is used.
<ntlmTokenReceivedPage>	This parameter is not required. Use to supply the URI of the resource with a response to be used when an NTLM token is received. If this parameter is not specified, the default NTLM token received error page is used.

The following is an example of the **modifySpnegoFilter** command:

```
wsadmin>$AdminTask modifySpnegoFilter {
  -hostName ks.austin.ibm.com
  -krb5Realm WSSEC.AUSTIN.IBM.COM}
```

Delete SPNEGO web authentication filter

Use the **deleteSpnegoFilter** command to remove SPNEGO a Web authentication filter from the security configuration. If a host name is not specified, all of the SPNEGO Web authentication filters are removed.

At the **wsadmin** prompt, enter the following command for help:

```
wsadmin>$AdminTask help deleteSpnegoFilter
```

You can use the following parameter with the **deleteSpnegoFilter** command:

Option	Description
<hostname>	This parameter is required. If the hostname is not specified, all of the SPNEGO Web authentication filters are deleted.

The following is an example of the **deleteSpnegoFilter** command:

```
wsadmin> $AdminTask deleteSpnegoFilter {-hostName ks.austin.ibm.com}
```

Show SPNEGO web authentication filter

Use the **showSpnegoFilter** command to display a SPNEGO Web authentication filter in the security configuration. If a host name is not specified, all of the SPNEGO filters are displayed.

At the **wsadmin** prompt, enter the following command for help:

```
wsadmin>$AdminTask help showSpnegoFilter
```

You can use the following parameter with the **showSpnegoFilter** command:

Option	Description
<hostname>	This parameter is optional. If a long host name is not specified, all of the SPNEGO Web authentication filters are displayed.

The following is an example of the **showSpnegoFilter** command:

```
wsadmin> $AdminTask showSpnegoFilter {-hostName ks.austin.ibm.com}
```

Kerberos authentication commands

Use **wsadmin** commands to create, modify or delete Kerberos as the authentication mechanism for WebSphere Application Server.

Create Kerberos authentication mechanism

Note:

The following items are required before you attempt to use the **createKrbAuthMechanism** command to create the KRB5 authentication mechanism security object field in the security configuration file:

- If you do not already have a Kerberos configuration file (**krb5.ini** or **krb5.conf**), use the **createkrbConfigFile** command task to create the Kerberos configuration file. Read about “Creating a Kerberos configuration file” on page 270 for more information.
- You must have a Kerberos keytab file (**krb5.keytab**) that contains a Kerberos service principal name (SPN), <service name>/<fully qualified hostname>@KerberosRealm, for each machine that run WebSphere application servers. The service name can be anything; the default value is WAS. For example, if you have two application server machines, host1.austin.ibm.com and host2.austin.ibm.com, the Kerberos keytab file must contain the <service name>/host1.austin.ibm.com and <service name>/host2.austin.ibm.com SPNs and their Kerberos keys.

Use the **createKrbAuthMechanism** command to create the KRB5 authentication mechanism security object field in the security configuration file.

At the **wsadmin** prompt, enter the following command:

```
$AdminTask help createKrbAuthMechanism
```

You can use the following parameters with the **createKrbAuthMechanism** command:

Option	Description
<krb5Realm>	This parameter is optional. It indicates the Kerberos realm name. If you do not specify this parameter, the default Kerberos realm in the Kerberos configuration file is used.
<krb5Config>	This parameter is required. It indicates the directory location and file name of the configuration (krb5.ini or krb5.conf) file.
<krb5Keytab>	This parameter is optional. It indicates the directory location and file name of the Kerberos keytab file. If you do not specify this parameter, the default keytab in the Kerberos configuration file is used.
<serviceName>	This parameter is required. It indicates the Kerberos service name. The default Kerberos service name is WAS.
<trimUserName>	This parameter is optional. It removes the suffix of the principal user name, starting from the “@” that precedes the Kerberos realm name. This parameter is optional. The default value is true.
<enabledGssCredDelegate>	This parameter is not required. Use to indicate whether to extract and place the client GSS delegation credential in the subject. The default value is true. Note: If this parameter is true, and the runtime cannot extract the GSS delegation credential, the runtime logs a warning message.
<allowKrbAuthForCsiInbound>	This parameter is optional. It enables Kerberos authentication mechanism for Common Secure Interoperability (CSI) inbound. The default value is true.
<allowKrbAuthForCsiOutbound>	This parameter is required. It enables Kerberos authentication mechanism for CSI outbound. The default value is true.

The following is an example of the **createKrbAuthMechanism** command:

```
wsadmin>$AdminTask createKrbAuthMechanism {
  -krb5Realm WSSEC.AUSTIN.IBM.COM
  -krb5Config C:\\WINNT\\krb5.ini
  -krb5Keytab C:\\WINNT\\krb5.keytab
  -serviceName WAS }
```

Modify Kerberos authentication mechanism

Use the **modifyKrbAuthMechanism** command to make changes to the KRB5 authentication mechanism security object field in the security configuration file.

At the **wsadmin** prompt, enter the following command:

```
wsadmin>$AdminTask help modifyKrbAuthMechanism
```

You can use the following parameters with the **modifyKrbAuthMechanism** command:

Option	Description
<krb5Realm>	This parameter is optional. It indicates the Kerberos realm name. If you do not specify this parameter, the default Kerberos realm in the Kerberos configuration file is used.

Option	Description
<krb5Config>	This parameter is required. It indicates the directory location and file name of the configuration (krb5.ini or krb5.conf) file.
<krb5Keytab>	This parameter is optional. It indicates the directory location and file name of the Kerberos keytab file. If you do not specify this parameter, the default keytab in the Kerberos configuration file is used.
<serviceName>	This parameter is required. It indicates the Kerberos service name. The default Kerberos service name is WAS.
<trimUserName>	This parameter is optional. It removes the suffix of the principal user name, starting from the "@" that precedes the Kerberos realm name. This parameter is optional. The default value is true.
<enabledGssCredDelegate>	This parameter is not required. Use to indicate whether to extract and place the client GSS delegation credential in the subject. The default value is true. Note: If this parameter is true, and the runtime cannot extract the GSS delegation credential, the runtime logs a warning message.
<allowKrbAuthForCsiInbound>	This parameter is optional. It enables Kerberos authentication mechanism for Common Secure Interoperability (CSI) inbound. The default value is true.
<allowKrbAuthForCsiOutbound>	This parameter is optional. It enables Kerberos authentication mechanism for CSI outbound. The default value is true.

The following is an example of the **modifyKrbAuthMechanism** command:

```
wsadmin>$AdminTask modifyKrbAuthMechanism {
  -krb5Realm WSSEC.AUSTIN.IBM.COM
  -krb5Config C:\\WINNT\\krb5.ini
  -krb5Keytab C:\\WINNT\\krb5.keytab
  -serviceName WAS }
```

Delete Kerberos authentication mechanism

Use the **deleteKrbAuthMechanism** command to remove the KRB5 authentication mechanism security object field in the security configuration file.

At the **wsadmin** prompt, enter the following command:

```
wsadmin>$AdminTask help deleteKrbAuthMechanism
```

The following is an example of the **deleteKrbAuthMechanism** command:

```
wsadmin>$AdminTask deleteKrbAuthMechanism
```

Set active authentication mechanism

Use the **setActiveAuthMechanism** command to set the active authentication mechanism attribute in the security configuration.

At the **wsadmin** prompt, enter the following command:

```
wsadmin>$AdminTask help setActiveAuthMechanism
```

You can use the following parameter with the **setActiveAuthMechanism** command:

Option	Description
<authMechanismType>	This parameter is not required. It indicates the authentication mechanism type. The default is KRB5.

The following is an example of the **setActiveAuthMechanism** command:

```
wsadmin> $AdminTask setActiveAuthMechanism {-authMechanismType KRB5 }
```

LTPA_LDAPSecurityOn and LTPA_LDAPSecurityOff command usage

Use the examples in this topic to enable and disable LTPA/LDAP security, based on single sign-on using the LDAP user registry.

Enabling LTPA/LDAP global security:

- Use help to find out what arguments you need to provide:

Using Jacl:

```
LTPA_LDAPSecurityOn help
```

Example output:

```
Syntax: LTPA_LDAPSecurityOn server user password port domain
```

Using Jython:

```
LTPA_LDAPSecurityOn('help')
```

Example output:

```
Syntax: LTPA_LDAPSecurityOn(server, user, password, port, domain)
```

- Issue the call with the arguments provided to turn on LTPA/LDAP security:

Using Jacl:

```
LTPA_LDAPSecurityOn ldapServer1 user1 password1 660 ibm.com
```

Using Jython:

```
LTPA_LDAPSecurityOn('ldapServer1', 'user1', 'password1', '660', 'ibm.com')
```

Example output:

PLEASE READ BELOW:

Done with LTPA/LDAP security turning on process, now you need to restart all the processes to make it affected. Then you can start using the client with SOAP or RMI connector.

- If you use the SOAP connector to connect to the server, you need to modify the `soap.client.props` file in your `app_server_root/properties` directory. Update as below for SOAP connector:

```
com.ibm.SOAP.securityEnabled=true  
com.ibm.SOAP.loginUserId=user1  
com.ibm.SOAP.loginPassword=password1
```

- the `sas.client.props` file in your `app_server_root/properties` directory. Update as below for RMI connector:

```
com.ibm.CORBA.loginUserId=user1  
com.ibm.CORBA.loginPassword=password1
```

Disabling LTPA/LDAP global security:

- Issue the following call to turn off LTPA/LDAP global security

Using Jacl:

```
LTPA_LDAPSecurityOff
```

Using Jython:

```
LTPA_LDAPSecurityOff()
```

Example output:

LTPA/LDAP security is off now but you need to restart all the processes to make it affected.

Chapter 11. Configuring denial of service protection for the proxy server

You can configure a pair of properties on your proxy server or DMZ Secure Proxy Server for IBM WebSphere Application Server to limit your risk against denial of service attacks involving the buffering of large HTTP payloads.

Before you begin

Denial of service protection for the proxy server or the DMZ Secure Proxy Server for IBM WebSphere Application Server is not done during the creation of these servers. The proxy server or the DMZ Secure Proxy Server for IBM WebSphere Application Server that will include denial of service protection must already exist before following these steps.

About this task

Note:

Protection is now included to guard against a type of security breach known as a denial of service attack. This type of attack can typically send more traffic to a network address than the data buffers were designed to accommodate. The proxy server and DMZ Secure Proxy Server for IBM WebSphere Application Server have several properties that can be configured to limit your risk against denial of service attacks involving the buffering of large HTTP payloads.

A denial of service attack is a malicious type of security breach to a computer system that does not usually result in the theft of information or other security loss. This type of attack can typically send more traffic to a network address than the data buffers were designed to accommodate resulting in memory exhaustion. HTTP allows for the body of a message to be sent to an HTTP server as an HTTP request or an HTTP response. The body can be sent to the HTTP server in a series of sequential network writes instead of being sent in one large network write. This process is known as Transfer-Encoding chunking. The maximum size of a Transfer-Encoding: chunked response body and a Transfer-Encoding: chunked request body can be set to determine how much data is buffered before a network write is performed.

1. Click **Servers>Proxy Servers>***proxy_server_name*.
2. Under Proxy Settings, expand **HTTP Proxy Server Settings** and click **Denial of service protection**
3. Set the appropriate buffer size in kilobytes for **Maximum request body buffer size**. Buffering can lead to better performance because it will decrease the number of network writes for large payloads. The size of the buffer must not be configured too high or memory exhaustion can occur. To determine the optimal values for your environment, gradually increase the size of the buffer until the proper balance has been achieved.
4. Set the appropriate buffer size in kilobytes for **Maximum response body buffer size**. The same precaution must be taken when setting Maximum response body buffer size as were indicated for Maximum request body buffer size.

Results

After these steps have been properly complete, denial of service protection will be in place for your proxy server or DMZ Secure Proxy Server for IBM WebSphere Application Server.

Denial of Service Protection

Use these settings to define denial of service protection for the proxy server. The buffer chunk sizes are used to provide protection against denial of service attacks. It is very important to tune these settings to balance the level of protection with the performance impacts that can be experienced.

To view this administrative console page, click **Servers** → **Proxy Servers** → **<server_name>** → **HTTP Proxy Server Settings** → **Denial of Service Protection**.

Maximum request body buffer chunk size

This field sets the maximum amount of the request body data that is buffered in memory before it is sent. The proxy server and secured proxy server can buffer multiple network reads of a Transfer-Encoding: chunked request body before performing a network write. This can improve performance because buffering the data will decrease the number of network writes required for large payloads. This field should be adjusted with caution. If the buffer size is set too high, memory exhaustion, which is a common intent of denial of service attacks, might be experienced. You will need to test different buffer sizes to find the optimal level of balance between protection and performance for your system. The default value of this field is 32 kilobytes.

Maximum response body buffer chunk size

This field sets the maximum amount of the response body data that is buffered in memory before it is sent. The proxy server and secured proxy server can buffer multiple network reads of a Transfer-Encoding: chunked response body before performing a network write. This can improve performance because buffering the data will decrease the number of network writes required for large payloads. This field should be adjusted with caution. If the buffer size is set too high, memory exhaustion, which is a common intent of denial of service attacks, might be experienced. You will need to test different buffer sizes to find the optimal level of balance between protection and performance for your system. The default value of this field is 32 kilobytes.

Chapter 12. Tuning the security properties for the DMZ Secure Proxy Server for IBM WebSphere Application Server

When creating a DMZ Secure Proxy Server for IBM WebSphere Application Server, default security levels of high, medium and low are available. In addition to the predefined configuration levels, you can modify the security settings for your DMZ Secure Proxy Server for IBM WebSphere Application Server. When you choose to customize the settings, a qualitative value of high, medium or low is still assigned to inform you of the overall security level of your DMZ Secure Proxy Server for IBM WebSphere Application Server.

Before you begin

A DMZ Secure Proxy Server for IBM WebSphere Application Server must be installed before these steps can be completed. The DMZ Secure Proxy Server for IBM WebSphere Application Server profile must be registered with the AdminAgent for this panel to be available.

About this task

Installing the DMZ Secure Proxy Server for IBM WebSphere Application Server in the DMZ rather than the secured zone presents new security challenges. The DMZ Secure Proxy Server for IBM WebSphere Application Server has been equipped with various capabilities to provide protection for meeting these challenges. In addition to the predefined security configurations for your DMZ Secure Proxy Server for IBM WebSphere Application Server you can also tune the settings to customize the protection.

1. Click **Servers** → **Proxy Servers** → **<secured_proxy_server_name>** → **Custom security settings** to open up the Proxy security settings panel.
2. Choose your administration option for your DMZ Secure Proxy Server for IBM WebSphere Application Server.
 - Local Administration - Security level: medium and high.
This option allows two different types of administration. Managing the DMZ Secure Proxy Server for IBM WebSphere Application Server entirely using the wsadmin tool and loading updated profiles imported from inside the cell using the wsadmin tool are both considered Local Administration.
 - Remote Administration- Security level: low
3. Choose your routing option for your DMZ Secure Proxy Server for IBM WebSphere Application Server.
 - Static routing - Security level: high
 - Dynamic routing - Security level: low and medium
4. Choose your startup permission option for your DMZ Secure Proxy Server for IBM WebSphere Application Server.
 - Run as an unprivileged user - Security level: medium and high
 - Run as privileged user - Security level: low
5. Optional: If Run as an unprivileged user is selected, enter the user name or the user group whose identify the server should assume after startup has completed.
6. Choose your custom error page policy option for your DMZ Secure Proxy Server for IBM WebSphere Application Server.
 - Local error page handling - Security level: high
If you choose to use local error page handling, you need to select which error responses should use custom error messages. Select Handle local errors for responses generated by the proxy server and select Handle remote errors for responses generated by the backend server. Both options may be selected to use custom error messages for local and remote errors. Manage your error code mappings to determine the custom error pages to be used for specific responses.
 - Remote error page handling - Security level: low and medium

If you choose to use remote error page handling to include custom errors, you need to select which error responses should be customized. Select Handle local errors for responses generated by the proxy server and select Handle remote errors for responses generated by the backend server. Both options may be selected to use custom error messages for both local and remote errors. Manage the headers that should be sent to the custom error application and what status codes are to be recognized as errors.

Results

You have finished customizing the security settings for your DMZ Secure Proxy Server for IBM WebSphere Application Server. A qualitative value of high, medium or low has been calculated based on the settings you have chosen to demonstrate the Current DMZ Security level.

DMZ Secure Proxy Server for IBM WebSphere Application Server start up user permissions

The overall security level of the DMZ Secure Proxy Server for IBM WebSphere Application Server can be hardened by reverting the server process to run as an unprivileged user after startup. Although the DMZ Secure Proxy Server for IBM WebSphere Application Server must be started as a privileged user, changing the server process to run as an unprivileged user provides additional protection for local operating resources.

Like the proxy server, the DMZ Secure Proxy Server for IBM WebSphere Application Server must start under a privileged user because it requires authorization to initialize privileged ports. Ports lower than 1024 are considered privileged ports. After these ports are initialized and access to the protected ports is no longer required, it is possible to change the user association of the DMZ Secure Proxy Server for IBM WebSphere Application Server process. Altering the server process to run using the privileges of a user or a group that does not have authority to access the local operation system resources adds a layer of protection to those resources. The firewall helps protect local operating system resources for the proxy server, but as the DMZ Secure Proxy Server for IBM WebSphere Application Server is installed in the DMZ, this type of protection becomes a higher priority. Although changing the user association of the server process for the DMZ Secure Proxy Server for IBM WebSphere Application Server is not required, continuing to run as a privileged user does not use the extra layer of protection for local operation resources that is provided when the server process is changed to run as an unprivileged user.

Run as unprivileged user	This is considered a high and medium security level setting.
Run as privileged user	This is considered a low security level setting.

DMZ Secure Proxy Server for IBM WebSphere Application Server routing considerations

This topic summarizes some of the security implications that must be considered when choosing how your DMZ Secure Proxy Server for IBM WebSphere Application Server will match incoming HTTP requests to an application or routing rule.

The DMZ Secure Proxy Server for IBM WebSphere Application Server can be configured to route requests either statically or dynamically. Using static routing specifies routing is performed using a flat configuration file using routing precedence that is inherent to the ordering of the directives. Requests can also be routed dynamically using a best match mechanism that determines the installed application or routing rule that corresponds to a specific request. The DMZ Secure Proxy Server for IBM WebSphere Application Server will dynamically discover the best route to a destination and distribute to servers with like protocols. It is considered more secure to use static routing than dynamic routing. This is because dynamic routing requires the secured proxy server to communicate with the application server to map requests where as

static routing is self contained by the secure proxy server requiring no additional communication.

Static Routing	This is considered a high security level setting.
Dynamic Routing	This is considered a low and medium security level setting.

DMZ Secure Proxy Server for IBM WebSphere Application Server administration options

The DMZ Secure Proxy Server for IBM WebSphere Application Server is administered differently than the WebSphere proxy server. The DMZ Secure Proxy Server for IBM WebSphere Application Server is a separate binary installed in the DMZ. Installing the DMZ Secure Proxy Server for IBM WebSphere Application Server in the DMZ requires that administration be managed differently for security reasons. Several administrative options are available for administering the DMZ Secure Proxy Server for IBM WebSphere Application Server to provide different levels of balance between security and usability.

The most secure way to administer the DMZ Secure Proxy Server for IBM WebSphere Application Server is locally using the wsadmin tool. The DMZ Secure Proxy Server for IBM WebSphere Application Server does not have web container therefore local administration can only be done via the command line. Using the wsadmin commands locally to manage the DMZ Secure Proxy Server for IBM WebSphere Application Server is the most secure option available because it does not require any external listening ports to be opened.

The DMZ Secure Proxy Server for IBM WebSphere Application Server configurations can also be managed within the network deployment application server cell and then imported locally using the wsadmin commands. The configurations are maintained inside the cell as configuration only profiles. The profiles are registered with the Admin Agent and are then managed using the administrative console. After you implement any changes to the profile, you export the configuration to an configuration archive (CAR) file using the exportProxyProfile or exportProxyServer wsadmin commands. After you transmit the CAR file to the local DMZ Secure Proxy Server for IBM WebSphere Application Server installation using ftp, the CAR file is imported using the importProxyProfile or importProxyServer wsadmin commands. This option is also considered to be local administration.

The DMZ Secure Proxy Server for IBM WebSphere Application Server can be managed remotely from the Job Manager console. You can create, delete, start, stop and modify the DMZ Secure Proxy Server for IBM WebSphere Application Server configuration using the Job Manager console. However, typically before managing the DMZ Secure Proxy Server for IBM WebSphere Application Server using the Job Manager console, the DMZ Secure Proxy Server for IBM WebSphere Application Servers would be created and configured using one of the other administrative options explained above. Using the Job Manager console is useful for centrally managing multiple DMZ Secure Proxy Server for IBM WebSphere Application Servers but is less secure than the other options because it requires the SOAP connector listener port to be opened for communication with the Job Manager console.

Error handling security considerations for the DMZ Secure Proxy Server for IBM WebSphere Application Server

The overall security level of the DMZ Secure Proxy Server for IBM WebSphere Application Server is partially determined by the choices made regarding the handling of custom errors.

You can define a custom error page for each error code or a group of error codes on errors generated by the proxy server or the application server. This is done using HTTP status codes in responses to generate uniform customized error pages for the application. For security reasons, you can ensure that the error pages are read from the local file system instead of being forwarded to a custom remote application. Choosing this option limits the code path and eliminates the need for a potentially unauthorized application

to be run as the error message is generated based on a flat file. For more information about the error handling for the secured proxy server, see Overview of the custom error page policy.

The following security level settings are used when evaluating a custom security level. Local error page handling is used for all of the predefined security levels.

Local error page handling	This is considered a high security level setting.
Remote error page handling	This is considered a medium and low security level setting.

Proxy security level properties

These settings describe the attributes and policies that define the security level of a secured proxy server. The overall security level of the secured proxy server is set to the weakest level of security assigned to any of the individual settings.

To view this administrative console page, click **Servers > Proxy Servers > server_name > Custom security settings**. This panel will only be available for a secure proxy server profile that has been registered with the AdminAgent.

Current security level

A qualitative security level based on an evaluation of the current security related configuration values.

The possible values for Current DMZ Security are high, medium, low. During creation of the secured proxy server, default configurations of high, medium and low are available. You are also able to customize these security settings resulting in the Current DMZ Security level being calculated by the system. Each custom setting has an assigned value of high, medium or low. The overall security level is equal to value of the setting that is considered the least secure. For example, to have an overall security level of high, all settings must be configured to the values associated with a high level of security. If any of the settings are configured with a less secure value, the overall security level is the value of that setting.

Administration

Table 59. Administration options

Option	Used as the default value in the predefined security levels	Description
Local administration	The default value for the Medium and the High security levels	Specifies that administration of the secure proxy server can only be performed using wsadmin commands performed locally on the system.
Remote administration	The default value for the Low security level	Specifies that remote administration of the secure proxy server is permitted.

Routing

Table 60. Routing options

Option	Used as the default value in the predefined security levels	Description
Static routing	The default value for the High security level	Specifies that the proxy server will make routing determinations from routing information based on flat files on the file system. This is for Hypertext Transfer Protocol (HTTP) only
Dynamic routing	The default value for the Low and the Medium security levels	Specifies that the proxy server will dynamically discover the best route to a destination and distribute to servers with like protocols.

Start-up permissions

Table 61. Start-up permission options

Option	Used as the default value in the predefined security levels	Description
Run as an unprivileged user	The default value for the Medium and the High security levels	Specifies that the server process will revert to a predefined unprivileged user after start-up has completed.
Run as a privileged user	The default value for the Low security level	Specifies that the server process does not revert to an unprivileged user after startup. It is a requirement that the proxy server start under a privileged user as it initializes privileged ports. Ports lower than 1024 are considered privileged ports. Under this setting, the effective user of the server process continues to be the privileged user. This setting does not provide additional hardening to the access of the server process to the local operation system resources. This is considered a low security level setting.

Custom Error Page Policy

Table 62. Error page options

Option	Used as the default value in the predefined security levels	Description
Local error page handling	The default value for the Low, the Medium and the High security levels	Specifies that error responses will be generated from flat custom error page files stored locally on the local file system.
Remote error page handling	None	Specifies to route error responses to a remote custom application deployed on a back-end server. This application will generate a custom response for the error

Local error page handling

- **Handle errors generated by the proxy server**

Specifies if errors generated by the proxy server should be handled with the custom static error pages stored on the local file system. If this is not selected then the default error messages will be used instead of any custom error pages.

- **Handle errors generated by application servers**

Specifies if errors generated by the backend server should be handled with the custom static error pages stored on the local file system. If this is not selected then the default error messages will be used instead of any custom error pages.

- **Error mappings**

Specifies the error codes to match with specific static error pages stored on the file system. You can use a relative file path under the configured static file document root to assign a custom error file to be used for a specific error code or group of error codes. The wildcard character, *, is used to assign error files to groups of error codes.

Remote error page handling

- **Error page generation application URI**

Specifies the URI for the custom error page generation application.

- **Handle errors generated by the proxy server**

Specifies if errors generated by the proxy server should be handled with the custom error application deployed on the application server. If this is not selected then the default error messages will be used instead of any custom error pages.

- **Handle errors generated by application servers**

Specifies if errors generated by the backend server should be handled with the custom error application deployed on the application server. If this is not selected then the default error messages will be used instead of any custom error pages.

- **Headers to forward to Error page Application**

Specifies a list of the headers from the original request to forward to the error page generation application.

- **HTTP status codes that are to be recognized as errors**

Specifies a list of the status codes in a response that should be directed to the error page generation application.

Chapter 13. Tuning, hardening, and maintaining

After installing WebSphere Application Server, there are several considerations for tuning, strengthening, and maintaining your security configuration.

About this task

The following topics are covered in this section:

- **Tuning security configurations** You can tune your security configuration to balance performance with function. You can achieve this balance following considerations for tuning general security, Common Secure Interoperability version 2 (CSIv2), Lightweight Directory Access Protocol (LDAP) authentication, Web authentication, and authorization. For more information on tuning security, see “Tuning security configurations.”
- **Hardening security configurations** Several methods exist that you can use to protect your infrastructure and applications from different forms of attack. For more information on hardening your security, see “Hardening security configurations” on page 1281.
- **Securing passwords in files** Password encryption and encoding can add protect to passwords existing in files. For more information on encoding and encrypting passwords, see “Securing passwords in files” on page 1281.

Tuning security configurations

You can tune security to balance performance with function. You can achieve this balance following considerations for tuning general security, Common Secure Interoperability version 2 (CSIv2), Lightweight Directory Access Protocol (LDAP) authentication, Web authentication, and authorization.

About this task

Performance issues typically involve trade-offs between function and speed. Usually, the more function and the more processing that are involved, the slower the performance. Consider what type of security is necessary and what you can disable in your environment. For example, if your application servers are running in a Virtual Private Network (VPN), consider whether you can disable Secure Sockets Layer (SSL). If you have a lot of users, can they be mapped to groups and then associated to your Java Platform, Enterprise Edition (Java EE) roles? These questions are things to consider when designing your security infrastructure.

- Consider the following recommendations for tuning general security.
 - Consider disabling Java 2 security manager if you know exactly what code is put onto your server and you do not need to protect process resources. Remember that in doing so, you put your local resources at some risk.
 - Consider propagating new security settings to all nodes before restarting the deployment manager and node agents, to change the new security policy.

If your security configurations are not consistent across all servers, you get access denied errors. Therefore, you must propagate new security settings when enabling or disabling administrative security.

Configuration changes are generally propagated using configuration synchronization. If auto-synchronization is enabled, you can wait for the automatic synchronization interval to pass, or you can force synchronization before the synchronization interval expires. If you are using manual synchronization, you must synchronize all the nodes.

If the cell is in a configuration state and the security policy is mixed with nodes that have security enabled and disabled, you can use the syncNode utility to synchronize the nodes where the new settings are not propagated.

For more detailed information about enabling security in a distributed environment, see “Enabling security for the realm” on page 53.

- Consider increasing the cache and token timeout if you feel your environment is secure enough. By increasing these values, you have to re-authenticate less often. This action supports subsequent requests to reuse the credentials that already are created. The downside of increasing the token timeout is the exposure of having a token hacked and providing the hacker more time to hack into the system before the token expires. You can use security cache properties to determine the initial size of the primary and secondary hashtable caches, which affect the frequency of rehashing and the distribution of the hash algorithms.

See the article “Authentication cache settings” on page 107 for a list of these properties.

- Consider changing your administrative connector from Simple Object Access Protocol (SOAP) to Remote Method Invocation (RMI) because RMI uses stateful connections while SOAP is completely stateless. Run a benchmark to determine if the performance is improved in your environment.
 - Use the wsadmin script to complete the access IDs for all the users and groups to speed up the application startup. Complete this action if applications contain many users or groups, or if applications are stopped and started frequently. WebSphere Application Server maps user and group names to unique access IDs in the authorization table. The exact format of the access ID depends on the repository. The access ID can only be determined during and after application deployment. Authorization tables created during assembly time do not have the proper access IDs. See Commands for the AdminApp object for more information about how to update access IDs.
 - Consider tuning the Object Request Broker (ORB) because it is a factor in enterprise bean performance with or without security enabled. Refer to the ORB tuning guidelines topic.
 - If using SSL, enable the SSL session tracking mechanism option as described in the article, Session management settings.
 - In some cases, using the unrestricted Java Cryptography Extension (JCE) policy file can improve performance. Refer to the article, Tuning Web services security.
 - Distributing the workload to multiple Java virtual machines (JVMs) instead of a single JVM on a single machine can improve the security performance because there is less contention for authorization decisions.
- Consider the following steps to tune Common Secure Interoperability version 2 (CSIv2).
 - Consider using Secure Sockets Layer (SSL) client certificates instead of a user ID and password to authenticate Java clients. Because you are already making the SSL connection, using mutual authentication adds little overhead while it removes the service context that contains the user ID and password completely.
 - If you send a large amount of data that is not very security sensitive, reduce the strength of your ciphers. The more data you have to bulk encrypt and the stronger the cipher, the longer this action takes. If the data is not sensitive, do not waste your processing with 128-bit ciphers.
 - Consider putting only an asterisk (*) in the trusted server ID list (meaning trust all servers) when you use identity assertion for downstream delegation. Use SSL mutual authentication between servers to provide this trust. Adding this extra step in the SSL handshake performs better than having to fully authenticate the upstream server and check the trusted list. When an asterisk (*) is used, the identity token is trusted. The SSL connection trusts the server through client certificate authentication.
 - Ensure that stateful sessions are enabled for CSIv2. This is the default, but requires authentication only on the first request and on any subsequent token expirations.
 - If you are communicating only with WebSphere Application Server Version 5 or higher servers, make the Active Authentication Protocol CSI, instead of CSI and SAS. This action removes an interceptor invocation for every request on both the client and server sides.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

- Consider the following steps to tune Lightweight Directory Access Protocol (LDAP) authentication.
 1. In the administration console, click **Security > Global security**.

2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry** and click **Configure**.
 3. Select the **Ignore case for authorization** option in the standalone LDAP registry configuration, when case-sensitivity is not important.
 4. Select the **Reuse connection** option.
 5. Use the cache features that your LDAP server supports.
 6. Choose either the IBM Tivoli Directory Server or SecureWay directory type, if you are using an IBM Tivoli Directory Server. The IBM Tivoli Directory Server yields improved performance because it is programmed to use the new group membership attributes to improve group membership searches. However, authorization must be case insensitive to use IBM Tivoli Directory Server.
 7. Choose either iPlanet Directory Server (also known as Sun ONE) or Netscape as the directory if you are an iPlanet Directory user. Using the iPlanet Directory Server directory can increase performance in group membership lookup. However, use **Role** only for group mechanisms.
- Consider the following steps to tune Web authentication.
 - Increase the cache and token timeout values if you feel your environment is secure enough. The Web authentication information is stored in these caches and as long as the authentication information is in the cache, the login module is not invoked to authenticate the user. This supports subsequent requests to reuse the credentials that are already created. A disadvantage of increasing the token timeout is the exposure of having a token stolen and providing the thief more time to hack into the system before the token expires.
 - Enable single sign-on (SSO). To configure SSO, click **Security > Global security**. Under Web security, click **Single sign-on (SSO)**.
SSO is only available when you configure **LTPA** as the authentication mechanism in the Authentication mechanisms and expiration panel. Although you can select Simple WebSphere Authentication Mechanism (SWAM) as the authentication mechanism on the Authentication mechanisms and expiration panel, SWAM is deprecated in Version 7.0 and does not support SSO. When you select SSO, a single authentication to one application server is enough to make requests to multiple application servers in the same SSO domain. Some situations exist where SSO is not a desirable and you do not want to use it in those situations.
 - Disable or enabling the **Web Inbound Security Attribute Propagation** option on the Single sign-on (SSO) panel if the function is not required. In some cases, having the function enabled can improve performance. This improvement is most likely for higher volume cases where a considerable number of user registry calls reduces performance. In other cases, having the feature disabled can improve performance. This improvement is most likely when the user registry calls do not take considerable resources.
 - The following two custom properties might help to improve performance when security attribute propagation is enabled:
 - **com.ibm.CSI.propagateFirstCallerOnly**
When this custom property is set to true the first caller in the propagation token that stays on the thread is logged when security attribute propagation is enabled. Without setting this property, all of the caller switches are logged, which can affect performance.
 - **com.ibm.CSI.disablePropagationCallerList**
When this custom property is set to true the ability to add a caller or host list in the propagation token is completely disabled. This function is beneficial when the caller or host list in the propagation token is not needed in the environment.
 - Consider the following steps to tune authorization.
 - Map your users to groups in the user registry. Associate the groups with your Java Platform, Enterprise Edition (Java EE) roles. This association greatly improves performance when the number of users increases.
 - Judiciously assign method-permissions for enterprise beans. For example, you can use an asterisk (*) to indicate all the methods in the method-name element. When all the methods in enterprise

beans require the same permission, use an asterisk (*) for the method-name to indicate all methods. This indication reduces the size of deployment descriptors and reduces the memory that is required to load the deployment descriptor. It also reduces the search time during method-permission match for the enterprise beans method.

- Judiciously assign security-constraints for servlets. For example, you can use the *.jsp URL pattern to apply the same authentication data constraints to indicate all JavaServer Pages (JSP) files. For a given URL, the exact match in the deployment descriptor takes precedence over the longest path match. Use the *.jsp, *.do, *.html extension match if no exact matches exist and longest path matches exist for a given URL in the security constraints.
- Use new tuning parameters when using Java 2 security. The new tuning parameters can improve performance significantly, and introduce a new concept called *Read-only Subject*, which enables a new cache for J2C Auth Subjects when using container-managed auth data aliases. If the J2C auth subject does not need to be modified after it is created, the following new tuning parameters can be used to improve Java 2 Security performance:
 - com.ibm.websphere.security.auth.j2c.cacheReadOnlyAuthDataSubjects=true
 - com.ibm.websphere.security.auth.j2c.readOnlyAuthDataSubjectCacheSize=50 (This is the maximum number of subjects in the hashtable of the cache. Once the cache reaches this size, some of the entries are purged. For better performance, this size should be equal to the number of unique subjects (cache based on uniqueness of user principal + auth data alias + managed connection factory instance) when role-based security and Java 2 security are used together).
- Use new tuning parameters to improve the performance of Security Attribute Propagation. The new tuning parameters can be set through custom properties in the administrative console to reduce the extra overhead of Security Attribute Propagation:
 - com.ibm.CSI.disablePropagationCallerList=true
 - com.ibm.CSI.propagateFirstCallerOnly=true (use if you want to track the first caller only).

Results

You always have a trade off between performance, feature, and security. Security typically adds more processing time to your requests, but for a good reason. Not all security features are required in your environment. When you decide to tune security, create a benchmark before making any change to ensure that the change is improving performance.

What to do next

In a large scale deployment, performance is very important. Running benchmark measurements with different combinations of features can help you to determine the best performance versus the benefit of configuration for your environment. Continue to run benchmarks if anything changes in your environment, to help determine the impact of these changes.

Secure Sockets Layer performance tips

Use this page to learn about Secure Sockets Layer (SSL) performance tips. Be sure to consider that performance issues typically involve trade-offs between function and speed. Usually, the more function and the more processing that are involved, the slower the performance.

The following are two types of Secure Sockets Layer (SSL) performance:

- Handshake
- Bulk encryption and decryption

When an SSL connection is established, an SSL handshake occurs. After a connection is made, SSL performs bulk encryption and decryption for each read-write. The performance cost of an SSL handshake is much larger than that of bulk encryption and decryption.

To enhance SSL performance, decrease the number of individual SSL connections and handshakes.

Decreasing the number of connections increases performance for secure communication through SSL connections, as well as non-secure communication through simple Transmission Control Protocol/Internet Protocol (TCP/IP) connections. One way to decrease individual SSL connections is to use a browser that supports HTTP 1.1. Decreasing individual SSL connections can be impossible if you cannot upgrade to HTTP 1.1.

Another common approach is to decrease the number of connections (both TCP/IP and SSL) between two WebSphere Application Server components. The following guidelines help to verify the HTTP transport of the application server is configured so that the Web server plug-in does not repeatedly reopen new connections to the application server:

- Verify that the maximum number of keep alives are, at minimum, as large as the maximum number of requests per thread of the Web server (or maximum number of processes for IBM HTTP Server on UNIX). Make sure that the Web server plug-in is capable of obtaining a keep alive connection for every possible concurrent connection to the application server. Otherwise, the application server closes the connection after a single request is processed. Also, the maximum number of threads in the Web container thread pool should be larger than the maximum number of keep alives, to prevent the keep alive connections from consuming the Web container threads.

Note: HTTP Transports have been deprecated. For instructions on how to set a maximum keep alive value for channel based configurations, see HTTP transport channel settings.

- Increase the maximum number of requests per keep alive connection. The default value is 100, which means the application server closes the connection from the plug-in after 100 requests. The plug-in then has to open a new connection. The purpose of this parameter is to prevent denial of service attacks when connecting to the application server and preventing continuous send requests to tie up threads in the application server.

- Use a hardware accelerator if the system performs several SSL handshakes.


Hardware accelerators currently supported by WebSphere Application Server only increase the SSL handshake performance, not the bulk encryption and decryption. An accelerator typically only benefits the Web server because Web server connections are short-lived. All other SSL connections in WebSphere Application Server are long-lived.

- Use an alternative cipher suite with better performance.

The performance of a cipher suite is different with software and hardware. Just because a cipher suite performs better in software does not mean a cipher suite will perform better with hardware. Some algorithms are typically inefficient in hardware, for example, Data Encryption Standard (DES) and triple-strength DES (3DES); however, specialized hardware can provide efficient implementations of these same algorithms.

The performance of bulk encryption and decryption is affected by the cipher suite used for an individual SSL connection. The following chart displays the performance of each cipher suite. The test software calculating the data was Java Secure Socket Extension (JSSE) for both the client and server software, which used no cryptographic hardware support. The test did not include the time to establish a connection, but only the time to transmit data through an established connection. Therefore, the data reveals the relative SSL performance of various cipher suites for long running connections.

Before establishing a connection, the client enables a single cipher suite for each test case. After the connection is established, the client times how long it takes to write an integer to the server and for the server to write the specified number of bytes back to the client. Varying the amount of data had negligible effects on the relative performance of the cipher suites.



An analysis of the above data reveals the following:

- Bulk encryption performance is only affected by what follows the WITH in the cipher suite name. This is expected since the portion before the WITH identifies the algorithm used only during the SSL handshake.
- MD5 and Secure Hash Algorithm (SHA) are the two hash algorithms used to provide data integrity. MD5 is generally faster than SHA, however, SHA is more secure than MD5.
- DES and RC2 are slower than RC4. Triple DES is the most secure, but the performance cost is high when using only software.
- The cipher suite providing the best performance while still providing privacy is SSL_RSA_WITH_RC4_128_MD5. Even though SSL_RSA_EXPORT_WITH_RC4_40_MD5 is cryptographically weaker than RSA_WITH_RC4_128_MD5, the performance for bulk encryption is the same. Therefore, as long as the SSL connection is a long-running connection, the difference in the performance of high and medium security levels is negligible. It is recommended that a security level of high be used, instead of medium, for all components participating in communication only among WebSphere Application Server products. Make sure that the connections are long running connections.

Tuning security

Use the following procedures to tune the performance, without compromising your security settings.

About this task

Enabling security decreases performance. The following tuning parameters provide ways to minimize this performance impact.

- Disable security on any application servers that do not need security. You can disable security in the administrative console by clicking **Security > Global security** and deselecting the **Enable administrative security** option.
- Fine-tune the **Authentication cache timeout** value on the Authentication mechanisms and expiration panel in the administrative console. For more information, see the “Global security settings” on page 55 topic.
- Configure the security cache properties. For more information, see the “Authentication cache settings” on page 107 topic.
- Enable the **Enable SSL ID tracking** option on the Session management panel in the administrative console. For more information, see the Session management settings topic.

- Improve the performance of Web services security by downloading a Java Cryptography Extension (JCE) unlimited jurisdiction policy file that does not have restrictions on cryptography strength. For more information, see the Tuning Web services security for Version 7.0 applications topic.
- Read the Secure Sockets Layer performance tips and “Tuning security configurations” on page 1275 topics for more information.

Hardening security configurations

There are several methods that you can use to protect the WebSphere Application Server infrastructure and applications from different forms of attack. Several different techniques can help with multiple forms of attack. Sometimes a single attack can leverage multiple forms of intrusion to achieve the end goal.

About this task

For example, in the simplest case, network sniffing can be used to obtain passwords and those passwords can then be used to mount an application-level attack. The following issues are discussed in IBM WebSphere Developer Technical Journal: WebSphere Application Server V5 advanced security and system hardening:

- Take preventative measures to protect the infrastructure.
- Make applications less vulnerable to attack.
- At a minimum, ensure administrative security is enabled in all WebSphere processes. This protects access to the administrative ConfigService interface and managed beans (MBeans) that enables control over the WebSphere process if it is compromised.
- Ensure Secure Sockets Layer (SSL) is used whenever possible, and mutual SSL whenever possible. However, mutual SSL requires all clients to supply a trusted personal certificate in order to connect.
- Remove any unnecessary certificate authority (CA) signer certificates from your trust stores.
- Change default keystore passwords during or after profile creation using the AdminTask `changeMultipleKeyStorePasswords` command.
- Change your Lightweight Third-Party Authentication (LTPA) keys periodically. By default, this occurs automatically every 12 weeks. If you want to disable this automatic regeneration, remember to manually generate a new set of keys on occasion.
- Common Secure Interoperability version 2 (CSIv2) inbound Basic authentication is supported in this release of WebSphere Application Server. This means that the authentication process is optional. Consider changing the authentication default to 'required'.

Securing passwords in files

Password encoding and encryption deters the casual observation of passwords in server configuration and property files.

About this task

The following topics can be used to add protection for passwords located in files:

- Encoding passwords in files WebSphere Application Server contains some encoded passwords that are not encrypted. The **PropFilePasswordEncoder** utility is included to encode these passwords. For more information on encoding passwords in a file, see “Encoding passwords in files” on page 1282.
- Enabling custom password encryption You need to protect passwords that are contained in your WebSphere Application Server configuration. You can added protection by creating a custom class for encrypting the passwords. For more information on custom password encryption, see “Enabling custom password encryption” on page 1285.

Encoding passwords in files

The purpose of password encoding is to deter casual observation of passwords in server configuration and property files. Use the **PropFilePasswordEncoder** utility to encode passwords stored in properties files. WebSphere Application Server does not provide a utility for decoding the passwords. Encoding is not sufficient to fully protect passwords. Native security is the primary mechanism for protecting passwords used in WebSphere Application Server configuration and property files.

About this task

WebSphere Application Server contains several encoded passwords in files that are not encrypted. WebSphere Application Server provides the **PropFilePasswordEncoder** utility, which you can use to encode passwords. The purpose of password encoding is to deter casual observation of passwords in server configuration and property files. The **PropFilePasswordEncoder** utility does not encode passwords that are contained within XML or XMI files. Instead, WebSphere Application Server automatically encodes the passwords in these files. XML and XMI files that contain encoded passwords include the following:

Table 63. XML and XMI files that contain encoded passwords

File name	Additional information
<code>profile_root/config/cells/cell_name/security.xml</code>	The following fields contain encoded passwords: <ul style="list-style-type: none"> • LTPA password • JAAS authentication data • User registry server password • LDAP user registry bind password • Keystore password • Truststore password • Cryptographic token device password
<code>war/WEB-INF/ibm_web_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>ejb_jar/META-INF/ibm_ejbjar_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>client_jar/META-INF/ibm_appclient_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>ear/META-INF/ibm_application_bnd.xml</code>	Specifies the passwords for the default basic authentication for the run as bindings within all the descriptors
<code>profile_root/config/cells/cell_name/nodes/node_name/servers/server_name/security.xml</code>	The following fields contain encoded passwords: <ul style="list-style-type: none"> • Keystore password • Truststore password • Cryptographic token device password • Session persistence password • DRS client data replication password
<code>profile_root/config/cells/cell_name/nodes/node_name/servers/server_name/resources.xml</code>	The following fields contain encoded passwords: <ul style="list-style-type: none"> • WAS40Datasource password • mailTransport password • mailStore password • MQQueue queue mgr password

Table 63. XML and XMI files that contain encoded passwords (continued)

File name	Additional information
<i>profile_root</i> /config/cells/ <i>cell_name</i> /ws-security.xml	
ibm-webservices-bnd.xmi	
ibm-webservicesclient-bnd.xmi	

You use the **PropFilePasswordEncoder** utility to encode the passwords in properties files. These files include:

Table 64. The PropFilePasswordEncoder utility - Partial File List

File name	Additional information
<i>profile_root</i> /properties/sas.client.props	Specifies the passwords for the following files: <ul style="list-style-type: none"> com.ibm.ssl.keyStorePassword com.ibm.ssl.trustStorePassword com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/sas.tools.properties	Specifies passwords for: <ul style="list-style-type: none"> com.ibm.ssl.keyStorePassword com.ibm.ssl.trustStorePassword com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/sas.stdclient.properties	Specifies passwords for: <ul style="list-style-type: none"> com.ibm.ssl.keyStorePassword com.ibm.ssl.trustStorePassword com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/wsserver.key	
<i>profile_root</i> /profiles/AppSrvXX/properties/sib.client.ssl.properties	Specifies passwords for: <ul style="list-style-type: none"> com.ibm.ssl.keyStorePassword com.ibm.ssl.trustStorePassword
<i>profile_root</i> /UDDIReg/scripts/UDDIUtilityTools.properties	Specifies passwords for: <ul style="list-style-type: none"> trustStore.password

To encode a password again in one of the previous files, complete the following steps:

1. Access the file using a text editor and type over the encoded password. The new password is shown is no longer encoded and must be re-encoded.
2. Use the PropFilePasswordEncoder.bat or the PropFilePasswordEncode.sh file in the *profile_root*/bin directory to encode the password again.

If you are encoding files that are not SAS properties files, type PropFilePasswordEncoder "*file_name*" *password_properties_list*

Note: When you use the **PropFilePasswordEncoder** utility, a prompt asks whether a backup version of the original file is required. If a backup version is required, a backup file (.bak), is created with the clear text password. Examine the results and then delete this backup file. It contains the unencrypted password. If you do not want to see this prompt, edit the PropFilePasswordEncoder utility and add the following Java system property as a parameter:
-Dcom.ibm.websphere.security.util.createBackup=true or
-Dcom.ibm.websphere.security.util.createBackup=false

A true value for the Java system property creates a backup file and a false value disables the backup file.

where:

"*file_name*" is the name of the z/SAS properties file, and *password_properties_list* is the name of the properties to encode within the file.

Note: Only the password should be encoded in this file using the **PropFilePasswordEncoder** tool. Use the **PropFilePasswordEncoder** utility to encode WebSphere Application Server password files only. The utility cannot encode passwords that are contained in XML files or other files that contain open and close tags.

Results

If you reopen the affected files, the passwords are encoded. WebSphere Application Server does not provide a utility for decoding the passwords.

PropFilePasswordEncoder command reference

The **PropFilePasswordEncoder** command encodes passwords that are located in plain text property files. This command encodes both Secure Authentication Server (SAS) property files and non-SAS property files. After you encode the passwords, a decoding command does not exist.

To encode passwords, you must run this command from the directory:

Syntax

The command syntax is as follows:

```
PropFilePasswordEncoder "file_name" { passwordPropertiesList | -SAS } { -noBackup | -Backup }  
  [ -profileName profile ] [ -help | -? ]
```

Parameters

The following option is available for the **PropFilePasswordEncoder** command:

file_name

This required parameter specifies the name of the file in which passwords are encoded.

passwordPropertiesList

This parameter is required if you are encoding passwords in property files other than the `sas.client.props` file. Specify one or more password properties that you want to encode. The password properties list should be delimited by commas.

-SAS

This parameter is required if you are encoding passwords in the `sas.client.props` file.

-noBackup

This parameter is optional and the default. The script does not create a backup file. The default value can be altered by adding following Java System Property:

```
"-Dcom.ibm.websphere.security.util.createBackup=true".
```

-Backup

This parameter is optional. The script creates a backup file, `<file_name>.bak`, which contains passwords in clear text.

-profileName

This parameter is optional. The profile value specifies an application server profile name. The script uses the password encoding algorithm that it retrieves from the specified profile. If you do not specify this parameter, the script uses the default profile.

-help or -?

If you specify this parameter, the script ignores all other parameters and displays usage text.

Enabling custom password encryption

You need to protect passwords that are contained in your WebSphere Application Server configuration. After creating your server profile, you can add protection by creating a custom class for encrypting the passwords.

Before you begin

Create your custom class for encrypting passwords. For more information, see “Plug point for custom password encryption” on page 835.

About this task

Complete the following steps to enable custom password encryption.

1. Add the following system properties for every server and client process. For server processes, update the `server.xml` file for each process. Add these properties as a `genericJvmArgument` argument preceded by a **-D** prefix.

```
com.ibm.wsspi.security.crypto.customPasswordEncryptionClass=  
    com.acme.myPasswordEncryptionClass  
com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=true
```

Note: If the custom encryption class name is `com.ibm.wsspi.security.crypto.CustomPasswordEncryptionImpl`, it is automatically enabled when this class is present in the classpath. Do not define the system properties that are listed previously when the custom implementation has this package and class name. To disable encryption for this class, you must specify `com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=false` as a system property.

2. Add the Java archive (JAR) file containing the implementation class to the `app_server_root/classes` directory so that the WebSphere Application Server runtime can load the file.
3. Restart all server processes.
4. Edit each configuration document that contains a password and save the configuration. All password fields are then run through the **WSEncoderDecoder** utility, which calls the plug point when it is enabled. The `{custom:alias}` tags are displayed in the configuration documents. The passwords, even though they are encrypted, are still Base64-encoded. They seem similar to encoded passwords, except for the tags difference.
5. Encrypt any passwords that are in client-side property files using the **PropsFilePasswordEncoder** (.bat or .sh) utility. This utility requires that the properties listed previously are defined as system properties in the script to encrypt new passwords instead of encoding them.
6. To decrypt passwords from client Java virtual machines (JVMs), add the properties listed previously as system properties for each client utility.
7. Ensure that all nodes have the custom encryption classes in their class paths prior to enabling this function. The order in which enablement occurs is important. When adding a new node to a cell that contains password encryption, the new node must contain the custom encryption classes prior to using the **addNode** command. Consider the following Network Deployment enablement scenarios:
 - a. The `StandAloneProfile` profile is encrypting passwords with a different key prior to federation to a deployment manager cell. For this scenario, you must uninstall custom password encryption to ensure that the configuration has `{xor}` tags preceding the passwords prior to running the **addNode** command. The same implementation of the plug point must be in the `/classes` directory prior to running the **addNode** command, and the proper configuration properties are set so that the new node can recognize the encrypted password format of the `security.xml` file after federation completes.
 - b. The `StandAloneProfile` profile does not have password encryption configured prior to federation to a deployment manager cell. The same implementation of the plug point must be in the `/classes`

directory prior to running the **addNode** command, and the proper configuration properties are set so that the new node can recognize the encrypted password format of the `security.xml` file after federation completes.

- c. If enabling custom password encryption in a cell with multiple nodes present, update the correct configuration properties and have the custom password encryption implementation class located on all nodes. Stop all processes in the cell, and then start the deployment manager. Use the administrative console to edit the security configuration and then save it. Verify that the passwords are encrypted by looking at the `security.xml` file to see if the passwords are preceded by `{custom:alias}` tags.
- d. Run the **syncNode** command on each node, and start each one individually. If any nodes fail to start, make sure that they have custom password encryption enabled properly in each `security.xml` file and that the implementation class is in the appropriate `/classes` directory for the platform.

Results

Custom password encryption is enabled.

What to do next

If custom password encryption fails or is no longer required, see “Disabling custom password encryption.”

Disabling custom password encryption

If custom password encryption fails or is no longer required, perform this task to disable custom password encryption.

Before you begin

Enable custom password encryption.

About this task

Complete the following steps to disable custom password encryption.

1. Change the `com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled` property to be `false` in the `security.xml` file, but leave the `com.ibm.wsspi.security.crypto.customPasswordEncryptionClass` property configured. Any passwords in the model that still have the `{custom:alias}` tag are decrypted by using the customer password encryption class.
2. If an encryption key is lost, any passwords that are encrypted with that key cannot be retrieved. To recover a password, retype the password in the password field in plaintext and save the document. The new password must be written out using encoding with the `{xor}` tag with scripting or from the administrative console.

```
com.ibm.wsspi.security.crypto.customPasswordEncryptionClass=  
    com.acme.myPasswordEncryptionClass  
com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=false
```

3. Restart all processes to make the changes effective.
4. Edit each configuration document that contains an encrypted password and save the configuration. All password fields are then run through the **WSEncoderDecoder** utility, which calls the plug point in the presence of the `{custom:alias}` tag. The `{xor}` tags display in the configuration documents again after the documents are saved.
5. Decrypt and encode any passwords that are in client-side property files using the **PropsFilePasswordEncoder** (.bat or .sh) utility. If the encryption class is specified, but custom encryption is disabled, running this utility converts the encryption to encoding and causes the `{xor}` tags to display again.

6. Disable custom password encryption from the client Java virtual machines (JVMs) by adding the system properties listed previously to all client scripts. This action enables the code to decrypt passwords, but this action is not used to encrypt them again. The {xor} algorithm becomes the default for encoding. Leave the custom password encryption class defined for a time in case any encrypted passwords still exist in the configuration.

Results

Custom password encryption is disabled.

Chapter 14. Troubleshooting security configurations

The following topics help to troubleshoot specific problems that are related to configuring and enabling security configurations.

About this task

Refer to Security components troubleshooting tips for instructions on how to troubleshoot errors that are related to security.

Refer to SPNEGO TAI troubleshooting tips for instructions on how to troubleshoot errors that are related to diagnosing Simple and Protected GSS-API Negotiation (SPNEGO) trust association interceptor (TAI) problems and exceptions.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

- Errors when configuring or enabling security
- Errors after enabling security
- Access problems after enabling security
- Errors after configuring or enabling Secure Sockets Layer
- Errors configuring Secure Sockets Layer encrypted access
- Single sign-on configuration troubleshooting tips
- Authorization provider troubleshooting tips

Security components troubleshooting tips

This document explains basic resources and steps for diagnosing security-related issues in WebSphere Application Server.

Basic resources and steps for diagnosing security-related issues in WebSphere Application Server include:

- What “Log files” on page 1290 to look at and what to look for in them.
- What to look at and what to look for “Using SDSF” on page 1291.
- “General approach for troubleshooting security-related issues” on page 1292 to isolating and resolving security problems.
- When and how to “Trace security” on page 1296.
- An overview and table of “CSIv2 CORBA minor codes” on page 1297.

The following security-related problems are addressed elsewhere in the information center:

- Errors and access problems after enabling security

After enabling security, a degradation in performance is realized. For more information about using unrestricted policy files, see the Enabling security for the realm section of the *Securing applications and their environment* PDF book.

- Errors after enabling SSL, or SSL-related error messages
- Errors trying to configure and enable security

If none of these steps solves the problem, check to see if the problem is 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 an overview of WebSphere Application Server security components such as Secure Authentication Services (SAS) and how they work in a distributed or an iSeries environment, refer to the *Securing applications and their environment* PDF book.

Note: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

Log files

When troubleshooting the security component, browse the Java Virtual Machine (JVM) logs for the server that hosts the resource you are trying to access. The following is a sample of messages you would expect to see from a server in which the security service has started successfully:

```
SASRas      A CWWSA0001I: Security configuration initialized.
SASRas      A CWWSA0002I: Authentication protocol: CSIV2/IBM
SASRas      A CWWSA0003I: Authentication mechanism: SWAM
SASRas      A CWWSA0004I: Principal name: MYHOSTNAME/aServerID
SASRas      A CWWSA0005I: SecurityCurrent registered.
SASRas      A CWWSA0006I: Security connection interceptor initialized.
SASRas      A CWWSA0007I: Client request interceptor registered.
SASRas      A CWWSA0008I: Server request interceptor registered.
SASRas      A CWWSA0009I: IOR interceptor registered.
NameServerImp I CWNMS0720I: Do Security service listener registration.
SecurityCompo A CWSCJ0242A: Security service is starting
UserRegistryI A CWSCJ0136I: Custom Registry:com.ibm.ws.security.registry.nt.
NTLocalDomainRegistryImpl has been initialized
SecurityCompo A CWSCJ0202A: Admin application initialized successfully
SecurityCompo A CWSCJ0203A: Naming application initialized successfully
SecurityCompo A CWSCJ0204A: Rolebased authorizer initialized successfully
SecurityCompo A CWSCJ0205A: Security Admin mBean registered successfully
SecurityCompo A CWSCJ0243A: Security service started successfully
SecurityCompo A CWSCJ0210A: Security enabled true
```

The following is an example of messages from a server which cannot start the security service, in this case because the administrative user ID and password given to communicate with the user registry is wrong, or the user registry itself is down or misconfigured:

```
SASRas      A CWWSA0001I: Security configuration initialized.
SASRas      A CWWSA0002I: Authentication protocol: CSIV2/IBM
SASRas      A CWWSA0003I: Authentication mechanism: SWAM
SASRas      A CWWSA0004I: Principal name: MYHOSTNAME/aServerID
SASRas      A CWWSA0005I: SecurityCurrent registered.
SASRas      A CWWSA0006I: Security connection interceptor initialized.
SASRas      A CWWSA0007I: Client request interceptor registered.
SASRas      A CWWSA0008I: Server request interceptor registered.
SASRas      A CWWSA0009I: IOR interceptor registered.
NameServerImp I CWNMS0720I: Do Security service listener registration.

SecurityCompo A CWSCJ0242A: Security service is starting
UserRegistryI A CWSCJ0136I: Custom Registry:com.ibm.ws.security.
registry.nt.NTLocalDomainRegistryImpl has been initialized
Authenticatio E CWSCJ4001E: Login failed for badID/<null>
javax.security.auth.login.LoginException: authentication failed: bad user/password
```

The following is an example of messages from a server for which Lightweight Directory Access Protocol (LDAP) has been specified as the security mechanism, but the LDAP keys have not been properly configured:

```

SASRas      A CWSA0001I: Security configuration initialized.
SASRas      A CWSA0002I: Authentication protocol: CSIV2/IBM
SASRas      A CWSA0003I: Authentication mechanism: LTPA
SASRas      A CWSA0004I: Principal name: MYHOSTNAME/anID
SASRas      A CWSA0005I: SecurityCurrent registered.
SASRas      A CWSA0006I: Security connection interceptor initialized.
SASRas      A CWSA0007I: Client request interceptor registered.
SASRas      A CWSA0008I: Server request interceptor registered.
SASRas      A CWSA0009I: IOR interceptor registered.
NameServerImp I CWNMS0720I: Do Security service listener registration.
SecurityCompo A CWSCJ0242A: Security service is starting
UserRegistryI A CWSCJ0136I: Custom Registry:com.ibm.ws.security.registry.nt.
NTLocalDomainRegistryImpl has been initialized
SecurityServe E CWSCJ0237E: One or more vital LTPAServerObject configuration
attributes are null or not available. The attributes and values are password :
LTPA password does exist, expiration time 30, private key <null>, public key <null>,
and shared key <null>.

```

A problem with the Secure Sockets Layer (SSL) configuration might lead to the following message. Ensure that the keystore location and keystore passwords are valid. Also, ensure the keystore has a valid personal certificate and that the personal certificate public key or certificate authority (CA) root has been extracted on put into the truststore.

```

SASRas      A CWSA0001I: Security configuration initialized.
SASRas      A CWSA0002I: Authentication protocol: CSIV2/IBM
SASRas      A CWSA0003I: Authentication mechanism: SWAM
SASRas      A CWSA0004I: Principal name: MYHOSTNAME/aServerId
SASRas      A CWSA0005I: SecurityCurrent registered.
SASRas      A CWSA0006I: Security connection interceptor initialized.
SASRas      A CWSA0007I: Client request interceptor registered.
SASRas      A CWSA0008I: Server request interceptor registered.
SASRas      A CWSA0009I: IOR interceptor registered.
SASRas      E CWSA0026E: [SecurityTaggedComponentAssistorImpl.register]
Exception connecting object to the ORB. Check the SSL configuration to ensure
that the SSL keyStore and trustStore properties are set properly. If the problem
persists, contact support for assistance. org.omg.CORBA.OBJ_ADAPTER:
ORB_CONNECT_ERROR (5) - couldn't get Server Subcontract minor code:
4942FB8F completed: No

```

Using SDSF

When troubleshooting the security component, use System Display and Search Facility (SDSF) to browse logs for the server that hosts the resource you are trying to access. The following sample of messages helps you see from a server in which the security service has started successfully:

```

+BBOM0001I com_ibm_authMechanisms_type_OID: No OID for this mechanism.
+BBOM0001I com_ibm_security_SAF_unauthenticated: WSGUEST.
+BBOM0001I com_ibm_security_SAF_EJBROLE_Audit_Messages_Suppress: 0.
+BBOM0001I com_ibm_ws_logging_zos_errorlog_format_cbe: NOT SET, 280
DEFAULT=0.
+BBOM0001I com_ibm_CSI_performClientAuthenticationRequired: 0.
+BBOM0001I com_ibm_CSI_performClientAuthenticationSupported: 1.
+BBOM0001I com_ibm_CSI_performTransportAssocSSLTLSRequired: 0.
+BBOM0001I com_ibm_CSI_performTransportAssocSSLTLSSupported: 1.
+BBOM0001I com_ibm_CSI_rmiInboundPropagationEnabled: 1.
+BBOM0001I com_ibm_CSI_rmiOutboundLoginEnabled: 0.
+BBOM0001I com_ibm_CSI_rmiOutboundPropagationEnabled: 1.
+BBOM0001I security_assertedID_IBM_accepted: 0.
+BBOM0001I security_assertedID_IBM_sent: 0.
+BBOM0001I security_disable_daemon_ssl: NOT SET, DEFAULT=0.
+BBOM0001I security_sslClientCerts_allowed: 0.
+BBOM0001I security_sslKeyring: NOT SET.
+BBOM0001I security_zOS_domainName: NOT SET.

```

```

+BB0M0001I security_zOS_domainType: 0.
+BB0M0001I security_zSAS_ssl_repertoire: SY1/DefaultIIOPSSL.
+BB0M0001I security_EnableRunAsIdentity: 0.
+BB0M0001I security_EnableSyncToOSthread: 0.
+BB0M0001I server_configured_system_name: SY1.
+BB0M0001I server_generic_short_name: BBOC001.
+BB0M0001I server_generic_uuid: 457
*** Message beginning with BB000222I apply to Java within ***
*** WebSphere Application Server Security ***
+BB000222I: SECJ6004I: Security Auditing is disabled.
+BB000222I: SECJ0215I: Successfully set JAAS login provider 631
configuration class to com.ibm.ws.security.auth.login.Configuration.
+BB000222I: SECJ0136I: Custom 632
Registry:com.ibm.ws.security.registry.zOS.SAFRegistryImpl has been initialized
+BB000222I: SECJ0157I: Loaded Vendor AuthorizationTable: 633
com.ibm.ws.security.core.SAFAuthorizationTableImpl

```

General approach for troubleshooting security-related issues

When troubleshooting security-related problems, the following questions are very helpful:

Does the problem occur when security is disabled?

This question is a good litmus test to determine that a problem is security related. However, just because a problem only occurs when security is enabled does not always make it a security problem. More troubleshooting is necessary to ensure the problem is really security-related.

Did security seem to initialize properly?

A lot of security code is visited during initialization. So you can see problems there first if the problem is configuration related.

The following sequence of messages that are generated in the SystemOut.log indicate normal code initialization of an application server. This sequence varies based on the configuration, but the messages are similar:

```

SASRas      A CWWSA0001I: Security configuration initialized.
SASRas      A CWWSA0002I: Authentication protocol: CSIV2/IBM
SASRas      A CWWSA0003I: Authentication mechanism: SWAM
SASRas      A CWWSA0004I: Principal name: BIRKT20/pbirk
SASRas      A CWWSA0005I: SecurityCurrent registered.
SASRas      A CWWSA0006I: Security connection interceptor initialized.
SASRas      A CWWSA0007I: Client request interceptor registered.
SASRas      A CWWSA0008I: Server request interceptor registered.
SASRas      A CWWSA0009I: IOR interceptor registered.
NameServerImp I CWNMS0720I: Do Security service listener registration.
SecurityCompo A CWSCJ0242A: Security service is starting
UserRegistryI A CWSCJ0136I: Custom Registry:com.ibm.ws.security.registry.nt.
NTLocalDomainRegistryImpl has been initialized
SecurityCompo A CWSCJ0202A: Admin application initialized successfully
SecurityCompo A CWSCJ0203A: Naming application initialized successfully
SecurityCompo A CWSCJ0204A: Rolebased authorizer initialized successfully
SecurityCompo A CWSCJ0205A: Security Admin mBean registered successfully
SecurityCompo A CWSCJ0243A: Security service started successfully

SecurityCompo A CWSCJ0210A: Security enabled true

```

The following sequence of messages generated in the SDSF active log indicate normal code initialization of an application server. Non-security messages have been removed from the sequence that follows. This sequence will vary based on the configuration, but the messages are similar:

```

Trace: 2005/05/06 17:27:31.539 01 t=8E96E0 c=UNK key=P8 (13007002)
  ThreadId: 0000000a
  FunctionName: printProperties
  SourceId: com.ibm.ws390.orb.CommonBridge
  Category: AUDIT
  ExtendedMessage: BBOJ0077I java.security.policy =
    /WebSphere/V6R1M0/AppServer/profiles/default/pr

```

Trace: 2005/05/06 17:27:31.779 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: printProperties
 SourceId: com.ibm.ws.390.orb.CommonBridge
 Category: AUDIT
 ExtendedMessage: BB0J0077I java.security.auth.login.config =
 /WebSphere/V6R1M0/AppServer/profiles/default/pr

Trace: 2005/05/06 17:27:40.892 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.core.SecurityDM
 SourceId: com.ibm.ws.security.core.SecurityDM
 Category: INFO
 ExtendedMessage: BB000222I: SECJ0231I: The Security component's FFDC
 Diagnostic Module com.ibm.ws.security.core.Secur
 red successfully: true.

Trace: 2005/05/06 17:27:40.892 01 t=8E96E0 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 932
 error message: BB000222I: SECJ0231I: The Security component's FFDC
 Diagnostic Module com.ibm.ws.security.core.Securit
 d successfully: true.

Trace: 2005/05/06 17:27:41.054 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.audit.AuditServiceImpl
 SourceId: com.ibm.ws.security.audit.AuditServiceImpl
 Category: AUDIT
 ExtendedMessage: BB000222I: SECJ6004I: Security Auditing is disabled.

Trace: 2005/05/06 17:27:41.282 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.core.distSecurityComponentImpl
 SourceId: com.ibm.ws.security.core.distSecurityComponentImpl
 Category: INFO
 ExtendedMessage: BB000222I: SECJ0309I: Java 2 Security is disabled.

Trace: 2005/05/06 17:27:41.282 01 t=8E96E0 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 932
 error message: BB000222I: SECJ0309I: Java 2 Security is disabled.

Trace: 2005/05/06 17:27:42.239 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.auth.login.Configuration
 SourceId: com.ibm.ws.security.auth.login.Configuration
 Category: AUDIT
 ExtendedMessage: BB000222I: SECJ0215I: Successfully set JAAS login
 provider configuration class to com.ibm.ws.securit
 Configuration.

Trace: 2005/05/06 17:27:42.253 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.core.distSecurityComponentImpl
 SourceId: com.ibm.ws.security.core.distSecurityComponentImpl
 Category: INFO
 ExtendedMessage: BB000222I: SECJ0212I: WCCM JAAS configuration information
 successfully pushed to login provider clas

Trace: 2005/05/06 17:27:42.254 01 t=8E96E0 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 932
 error message: BB000222I: SECJ0212I: WCCM JAAS configuration information
 successfully pushed to login provider class.

Trace: 2005/05/06 17:27:42.306 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.security.core.distSecurityComponentImpl
 SourceId: com.ibm.ws.security.core.distSecurityComponentImpl
 Category: INFO
 ExtendedMessage: BB000222I: SECJ0240I: Security service initialization
 completed successfully

Trace: 2005/05/06 17:27:42.306 01 t=8E96E0 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 932
 error message: BB000222I: SECJ0240I: Security service initialization
 completed successfully

Trace: 2005/05/06 17:27:42.952 01 t=8E96E0 c=UNK key=P8 (13007002)
 ThreadId: 0000000a
 FunctionName: com.ibm.ws.objectpool.ObjectPoolService
 SourceId: com.ibm.ws.objectpool.ObjectPoolService
 Category: INFO

```

ExtendedMessage: BB000222I: OBPL0007I: Object Pool Manager service
is disabled.
Trace: 2005/05/06 17:27:53.512 01 t=8E96E0 c=UNK key=P8 (13007002)
ThreadId: 0000000a
FunctionName: com.ibm.ws.security.registry.UserRegistryImpl
SourceId: com.ibm.ws.security.registry.UserRegistryImpl
Category: AUDIT
ExtendedMessage: BB000222I: SECJ0136I: Custom
Registry:com.ibm.ws.security.registry.zOS.SAFRegistryImpl
has been init
Trace: 2005/05/06 17:27:55.229 01 t=8E96E0 c=UNK key=P8 (13007002)
ThreadId: 0000000a
FunctionName: com.ibm.ws.security.role.PluggableAuthorizationTableProxy
SourceId: com.ibm.ws.security.role.PluggableAuthorizationTableProxy
Category: AUDIT
ExtendedMessage: BB000222I: SECJ0157I: Loaded Vendor
AuthorizationTable: com.ibm.ws.security.core.SAFAuthorizationTab
Trace: 2005/05/06 17:27:56.481 01 t=8E96E0 c=UNK key=P8 (13007002)
ThreadId: 0000000a
FunctionName: com.ibm.ws.security.core.distSecurityComponentImpl
SourceId: com.ibm.ws.security.core.distSecurityComponentImpl
Category: INFO
ExtendedMessage: BB000222I: SECJ0243I: Security service started successfully
Trace: 2005/05/06 17:27:56.481 01 t=8E96E0 c=UNK key=P8 (0000000A)
Description: Log Boss/390 Error
from filename: ./bborjtr.cpp
at line: 932
error message: BB000222I: SECJ0243I: Security service started successfully
Trace: 2005/05/06 17:27:56.482 01 t=8E96E0 c=UNK key=P8 (13007002)
ThreadId: 0000000a
FunctionName: com.ibm.ws.security.core.distSecurityComponentImpl
SourceId: com.ibm.ws.security.core.distSecurityComponentImpl
Category: INFO
ExtendedMessage: BB000222I: SECJ0210I: Security enabled true
Trace: 2005/05/06 17:27:56.483 01 t=8E96E0 c=UNK key=P8 (0000000A)
Description: Log Boss/390 Error
from filename: ./bborjtr.cpp
at line: 932
error message: BB000222I: SECJ0210I: Security enabled true

```

Is there a stack trace or exception printed in the system log file?

A single stack trace tells a lot about the problem. What code initiated the code that failed? What is the failing component? Which class did the failure actually come from? Sometimes the stack trace is all that is needed to solve the problem and it can pinpoint the root cause. Other times, it can only give us a clue, and can actually be misleading. When support analyzes a stack trace, they can request additional trace if it is not clear what the problem is. If it seems to be security-related and the solution cannot be determined from the stack trace or problem description, you are asked to gather the following trace specification: SASRas=all=enabled:com.ibm.ws.security.*=all=enabled from all processes involved.

Is this a distributed security problem or a local security problem?

- If the problem is local, that is the code involved does not make a remote method invocation, then troubleshooting is isolated to a single process. It is important to know when a problem is local versus distributed because the behavior of the object request broker (ORB), among other components, is different between the two. When a remote method invocation takes place, an entirely different security code path is entered.
- When you know that the problem involves two or more servers, the techniques of troubleshooting change. You need to trace all the servers involved simultaneously so that the trace shows the client and server sides of the problem. Make sure the timestamps on all machines match as closely as possible so that you can find the request and reply pair from two different processes. Enable both Secure Authentication Services (SAS) or z/SAS and Security trace using the trace specification: SASRas=all=enabled:com.ibm.ws.security.*=all=enabled.

For more information on enabling trace, see [Enabling trace](#).

For more information on enabling trace, see [Working with Trace](#).

Is the problem related to authentication or authorization?

Most security problems fall under one of these two categories. Authentication is the process of determining who the caller is. Authorization is the process of validating that the caller has the proper authority to invoke the requested method. When authentication fails, typically this failure is

related to either the authentication protocol, authentication mechanism or user registry. When authorization fails, this is usually related to the application bindings from assembly and deployment and to the caller's identity who is accessing the method and the roles that are required by the method.

Is this a Web or EJB request?

Web requests have a completely different code path than Enterprise JavaBeans (EJB) requests. Different security features exist for Web requests than for EJB requests, requiring a completely different body of knowledge to resolve. For example, when using the Lightweight Third-Party Authentication (LTPA) authentication mechanism, the single sign-on feature (SSO) is available for Web requests but not for EJB requests. Web requests involve HTTP header information that is not required by EJB requests due to the protocol differences. Also, the Web container or servlet engine is involved in the entire process. Any of these components can be involved in the problem and all require consideration during troubleshooting, based on the type of request and where the failure occurs.

Secure EJB requests heavily involve the ORB and Naming components since they flow over the RMI/IIOP protocol. In addition, when Workload Manager (WLM) is enabled, other behavior changes in the code can be observed. All of these components interact closely for security to work properly in this environment. At times, trace in any or all of these components might be necessary to troubleshoot problems in this area.

The trace specification to begin with is `SASRas=all=enabled:com.ibm.ws.security.*=all=enabled`. ORB trace is also very beneficial when the SAS/Security trace does not seem to pinpoint the problem.

Does the problem seem to be related to the Secure Sockets Layer (SSL)?

SSL is a totally distinct separate layer of security. Troubleshooting SSL problems is usually separate from troubleshooting authentication and authorization problems, and you have many considerations. Usually, SSL problems are first-time setup problems because the configuration can be difficult. Each client must contain the signer certificate of the server. During mutual authentication, each server must contain the client's signer certificate. Also, there can be protocol differences (SSLv3 vs. Transport Layer Security (TLS)), and listener port problems related to stale Interoperable Object References (IORs), that is IORs from a server, that reflect the port prior to the server restarting.

For SSL problems, sometimes you get a request for an SSL trace to determine what is happening with the SSL handshake. The SSL handshake is the process that occurs when a client opens a socket to a server. If anything goes wrong with the key exchange, cipher exchange, and so on, the handshake fails and the socket is not valid. Tracing JSSE (the SSL implementation that is used in WebSphere Application Server) involves the following steps:

- Set the following system property on the client and server processes: `-Djavax.net.debug=true`. For the server, add the system property to the generic JVM arguments property of the JVM settings page. For more information on this task, refer to Java virtual machine settings section of the *Administering applications and their environment* PDF book.
- Turn on ORB trace as well.
- Recreate the problem.

The `SystemOut.log` of both processes contain the JSSE trace. You can find trace similar to the following example:

```
SSLConnection: install <com.ibm.sslite.e@3ae78375>
>> handleHandshakeV2 <com.ibm.sslite.e@3ae78375>
>> handshakeV2 type = 1
>> clientHello: SSLv2.
SSL client version: 3.0
...
...
JSSEContext: handleSession[Socket[addr=null,port=0,localport=0]]

<< sendServerHello.
SSL version: 3.0
SSL_RSA_WITH_RC4_128_MD5
```

```

HelloRandom
...
...
...
<< sendCertificate.
<< sendServerHelloDone.
>> handleData <com.ibm.sslite.e@3ae78375>
>> handleHandshake <com.ibm.sslite.e@3ae78375>
>> handshakeV3 type = 16

>> clientKeyExchange.
>> handleData <com.ibm.sslite.e@3ae78375>
>> handleChangeCipherSpec <com.ibm.sslite.e@3ae78375>
>> handleData <com.ibm.sslite.e@3ae78375>
>> handleHandshake <com.ibm.sslite.e@3ae78375>
>> handshakeV3 type = 20
>> finished.
<< sendChangeCipherSpec.
<< sendFinished.

```

Trace security

The classes that implement WebSphere Application Server security are:

- com.ibm.ws.security.*
- com.ibm.websphere.security.*
- com.ibm.WebSphereSecurityImpl.*
- SASRas
- com.ibm.ws.wim.* for tracing with a Virtual Member Manager (VMM) repository

To view detailed information on the run time behavior of security, enable trace on the following components and review the output:

- com.ibm.ws.security.*=all=enabled:com.ibm.WebSphereSecurityImpl.*=all=enabled:com.ibm.websphere.security.*=all=enabled. This trace statement collects the trace for the security runtime.
- com.ibm.ws.console.security.*=all=enabled. This trace statement collects the trace for the security center administrative console.
- SASRas=all=enabled. This trace statement collects the trace for SAS (low-level authentication logic).
- com.ibm.ws.wim.*=all=enabled:com.ibm.websphere.wim.*=all=enabled. This trace statement collects the trace for VMM.

Fine tuning SAS traces:

If a subset of classes need to be traced for the SAS/CSlv2 component, a system property can be specified with the class names comma separated:

```
com.ibm.CORBA.securityTraceFilter=SecurityConnectionInterceptorImpl, VaultImpl, ...
```

Fine tuning Security traces:

If a subset of packages need to be traced, specify a trace specification more detailed than com.ibm.ws.security.*=all=enabled. For example, to trace just dynamic policy code, you can specify com.ibm.ws.security.policy.*=all=enabled. To disable dynamic policy trace, you can specify com.ibm.ws.security.policy.*=all=disabled.

Configuring CSlv2, or SAS Trace Settings

Situations arise where reviewing trace for the CSlv2 or SAS authentication protocols can assist in troubleshooting difficult problems. This section describes how to enable to CSlv2 and SAS trace.

Enabling Client-Side CSlv2 and SAS Trace

To enable CSlv2 and SAS trace on a pure client, the following steps need to be taken:

- Edit the file TraceSettings.properties in the /WebSphere/AppServer/properties directory.
- In this file, change traceFileName= to point to the path in which you want the output file created. Make sure you put a double backslash (\\) between each subdirectory. For example, traceFileName=c:\\WebSphere\\AppServer\\logs\\sas_client.log
- In this file, add the trace specification string: SASRas=all=enabled. Any additional trace strings can be added on separate lines.

- Point to this file from within your client application. On the Java command line where you launch the client, add the following system property:
-DtraceSettingsFile=TraceSettings.properties.

Note: Do not give the fully qualified path to the TraceSettings.properties file. Make sure that the TraceSettings.properties file is in your class path.

Enabling Server-Side CSiv2 and SAS Trace

To enable SAS trace in an application server, complete the following:

- Add the trace specification, SASRas=all=enabled, to the server.xml file or add it to the Trace settings within the WebConsole GUI.
- Typically it is best to also trace the authorization security runtime in addition to the authentication protocol runtime. To do this, use the following two trace specifications in combination: SASRas=all=enabled:com.ibm.ws.security.*=all=enabled.
- When troubleshooting a connection type problem, it is beneficial to trace both CSiv2 and SAS or CSiv2 and z/SAS and the ORB. To do this, use the following three trace specifications:
SASRas=all=enabled:com.ibm.ws.security.*=all=enabled:ORBRas=all=enabled.
- In addition to adding these trace specifications, for ORB trace there are a couple of system properties that also need to be set. Go to the ORB settings in the GUI and add the following two properties: com.ibm.CORBA.Debug=true and com.ibm.CORBA.CommTrace=true.

CSiv2 CORBA minor codes

Whenever exceptions occur within the security code on either the client or server, the eventual exception becomes a Common Object Request Broker Architecture (CORBA) exception. Any exception that occurs gets embedded in a CORBA exception because the CORBA architecture is used by the security service for its own inter-process communication. CORBA exceptions are generic and indicate a problem in communication between two components. CORBA minor codes are more specific and indicate the underlying reason that a component could not complete a request.

The following shows the CORBA minor codes that a client can expect to receive after running a security-related request such as authentication. It also includes the CORBA exception type that the minor code appears in.

The following exception shows an example of a CORBA exception where the minor code is 49424300 and indicates Authentication Failure. Typically, a descriptive message is also included in the exception to assist in troubleshooting the problem. Here, the detailed message is: "Exception caught invoking authenticateBasicAuthData from SecurityServer for user jdoe. Reason: com.ibm.WebSphereSecurity.AuthenticationFailedException" which indicates that the authentication failed for user jdoe.

The completed field in the exception indicates whether the method was completed or not. In the case of a NO_PERMISSION, never invoke the message; therefore it is always completed:No. Other exceptions that are caught on the server side can have a completed status of "Maybe" or "Yes".

```
org.omg.CORBA.NO_PERMISSION: Caught WSSecurityContextException in
WSSecurityContext.acceptSecContext(),
reason: Major Code[0] Minor Code[0] Message[Exception caught invoking
authenticateBasicAuthData from SecurityServer for user jdoe. Reason:
com.ibm.WebSphereSecurity.AuthenticationFailedException] minor code: 49424300
completed: No
```

```
at com.ibm.ISecurityLocalObjectBaseL13Impl.PrincipalAuthFailReason.
map_auth_fail_to_minor_code(PrincipalAuthFailReason.java:83)
at com.ibm.ISecurityLocalObjectBaseL13Impl.CSIServerRI.receive_request
(CSIServerRI.java:1569)
at com.ibm.rmi.pi.InterceptorManager.iterateReceiveRequest
```



```

(InterceptorManager.java:739)
at com.ibm.CORBA.iiop.ServerDelegate.dispatch(ServerDelegate.java:398)
at com.ibm.rmi.iiop.ORB.process(ORB.java:313)
at com.ibm.CORBA.iiop.ORB.process(ORB.java:1581)
at com.ibm.rmi.iiop.GIOPConnection.doWork(GIOPConnection.java:1827)
at com.ibm.rmi.iiop.WorkUnitImpl.doWork(WorkUnitImpl.java:81)
at com.ibm.ejs.oa.pool.PooledThread.run(ThreadPool.java:91)
at com.ibm.ws.util.CachedThread.run(ThreadPool.java:149)

```

The following table shows the CORBA minor codes which a client can expect to receive after running a security-related request such as authentication. It also includes the CORBA exception type that the minor code would appear in.

Table 65.

Minor code name	Minor code value (in hex)	Exception type (all in the package of org.omg.CORBA.*)	Minor code description	Retry performed (when authenticationRetryEnabled = true)
AuthenticationFailed	49424300	NO_PERMISSION	This code is a generic authentication failed error. It does not give any details about whether or not the user ID or password is valid. Some user registries can choose to use this type of error code, others can choose to use the next three types that are more specific.	Yes
InterceptLocateException	494210B8	INTERNAL	This indicates a problem when processing an incoming locate request.	No
InvalidUserId	49424301	NO_PERMISSION	This code occurs when the registry returns bad user ID.	Yes
InvalidPassword	49424302	NO_PERMISSION	This code occurs when the registry returns a bad password.	Yes
InvalidSecurityCredentials	49424303	NO_PERMISSION	This is a generic error indicating that the credentials are bad for some reason. It might be that the right attributes are not set.	Yes, if client has BasicAuth credential (token based credential was rejected in the first place).
InvalidRealm	49424304	NO_PERMISSION	This code occurs when the REALM in the token received from the client does not match the server's current realm.	No
ValidationFailed	49424305	NO_PERMISSION	A validation failure occurs when a token is sent from the client or server to a target server but the token format or the expiration is not valid.	Yes, if client has BasicAuth credential (token based credential was rejected in the first place).
CredentialTokenExpired	49424306	NO_PERMISSION	This code is more specific about why the validation failed. In this case, the token has an absolute lifetime and the lifetime has expired. Therefore, it is no longer a valid token and cannot be used.	Yes, if client has BasicAuth credential (token based credential was rejected in the first place).
InvalidCredentialToken	49424307	NO_PERMISSION	This is more specific about why the validation failed. In this case, the token cannot be decrypted or the data within the token is not readable.	Yes, if client has BasicAuth credential (token based credential was rejected in the first place).
SessionDoesNotExist	49424308	NO_PERMISSION	This indicates that the CSIv2 session does not exist on the server. Typically, a retry occurs automatically and successfully creates a new session.	Yes

Table 65. (continued)

Minor code name	Minor code value (in hex)	Exception type (all in the package of org.omg.CORBA.*)	Minor code description	Retry performed (when authenticationRetryEnabled = true)
SessionConflictingEvidence	49424309	NO_PERMISSION	This indicates that a session already exists on the server that matches the context_id sent over by the client. However, the information provided by the client for this EstablishContext message is different from the information originally provided to establish the session.	Yes
SessionRejected	4942430A	NO_PERMISSION	This indicates that the session referenced by the client has been previously rejected by the server.	Yes
SecurityServerNotAvailable	4942430B	NO_PERMISSION	This error occurs when the server cannot contact the local or remote security server in order to authenticate or validate.	No
InvalidIdentityToken	4942430C	NO_PERMISSION	This error indicates that identity cannot be obtained from the identity token when Identity Assertion is enabled.	No
IdentityServerNotTrusted	4942430D	NO_PERMISSION	This indicates that the server ID of the sending server is not on the target server's trusted principal list.	No
InvalidMessage	4942430E	NO_PERMISSION	This indicates that the CSIV2 message format is not valid for the receiving server.	No
AuthenticationNotSupported	49421090	NO_PERMISSION	This error occurs when a mechanism does not support authentication (very rare).	No
InvalidSecurityMechanism	49421091	NO_PERMISSION	This is used to indicate that the specified security mechanism is not known.	No
CredentialNotAvailable	49421092	NO_PERMISSION	This indicates a credential is not available when it is required.	No
SecurityMechanismNotSupported	49421093	NO_PERMISSION	This error occurs when a security mechanism that is specified in the CSIV2 token is not implemented on the server.	No
ValidationNotSupported	49421094	NO_PERMISSION	This error occurs when a mechanism does not support validation, such as LocalOS. This error does not occur since the LocalOS credential is not a forwardable credential, therefore, validation never needs to be called on this credential.	No
CredentialTokenNotSet	49421095	NO_PERMISSION	This is used to indicate that the token inside the credential is null.	No
ServerConnectionFailed	494210A0	COMM_FAILURE	This error is used when a connection attempt fails.	Yes (via ORB retry)
CorbaSystemException	494210B0	INTERNAL	This code is a generic CORBA specific exception in system code.	No
JavaException	494210B1	INTERNAL	This is a generic error that indicated that an unexpected Java exception occurred.	No

Table 65. (continued)

Minor code name	Minor code value (in hex)	Exception type (all in the package of org.omg.CORBA.*)	Minor code description	Retry performed (when authenticationRetryEnabled = true)
ValuesIsNull	494210B2	INTERNAL	This code is used to indicate that a value or parameter that passed in is null.	No
EffectivePolicyNotPresent	494210B3	INTERNAL	This indicates that an effective policy object for CSiv2 is not present. This object is used to determine what security configuration features are specified.	No
NullPointerException	494210B4	INTERNAL	This code is used to indicate that a NullPointerException is caught in the runtime.	No
ErrorGettingClassInstance	494210B5	INTERNAL	This indicates a problem loading a class dynamically.	No
MalFormedParameters	494210B6	INTERNAL	This indicates parameters are not valid.	No
DuplicateSecurityAttributeType	494210B7	INTERNAL	This indicates a duplicate credential attribute that is specified during the set_attributes operation.	No
MethodNotImplemented	494210C0	NO_IMPLEMENT	This indicates that a method invoked is not implemented.	No
GSSFormatError	494210C5	BAD_PARAM	This code indicates that a Generic Security Services (GSS) encoding or decoding routine has created an exception.	No
TagComponentFormatError	494210C6	BAD_PARAM	This code indicates that a tag component cannot be read properly.	No
InvalidSecurityAttributeType	494210C7	BAD_PARAM	This code indicates an attribute type specified during the set_attributes operation is not a valid type.	No
SecurityConfigError	494210CA	INITIALIZE	This code indicates a problem exists between the client and server configuration.	No

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

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

Security configuration and enablement errors

Use this information to troubleshoot problems with configuring or enabling security.

What kind of error are you seeing?

- ““LTPA password not set. validation failed” message displayed as error in the administrative console after saving administrative or application security settings ” on page 1301
- “The setupClient.bat or setupClient.sh file is not working correctly” on page 1301
- **HP-UX** “Java HotSpot Server VM warning: Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a message occurs in the native_stdout.log file when enabling security on the HP-UX11i platform” on page 1301
- “WebSphere Application Server Version 6 is not working correctly with Enterprise Workload Manager (EWLM)” on page 1302

- If you successfully configured security, but are now having problems accessing Web resources or the administrative console, refer to Errors or access problems after enabling security.
- “NMSV0610I: A NamingException is being thrown from a javax.naming.Context implementation” on page 1302
- “When administrative security is enabled but application security is not enabled, the performance servlet displays authorization errors and cannot provide statistics” on page 1302
- ““Name value is invalid” displays when migrating users and groups after the JACC provider for Tivoli is configured” on page 1303
- “A Sun JDK can not read a PKCS12 keystore created by the Application Server” on page 1303

For general tips on diagnosing and resolving security-related problems, see the topic Troubleshooting the security component.

“LTPA password not set. validation failed” message displayed as error in the administrative console after saving administrative or application security settings

This error can be caused if, when configuring WebSphere Application Server security, LTPA is selected as the authentication mechanism and the LTPA password field is not set. To resolve this problem:

- Select **Security > Global security > Authentication mechanisms and expiration** .
- Complete the password and confirm password fields.
- Click **OK**.
- Try setting administrative or application security again.

The setupClient.bat or setupClient.sh file is not working correctly

The setupClient.bat file on Windows operating systems and the setupClient.sh file on Linux and UNIX-based platforms incorrectly specify the location of the SOAP security properties file.

Windows In the setupClient.bat file, the correct location is:

```
set CLIENTSOAP=-Dcom.ibm.SOAP.ConfigURL=file:%WAS_HOME%/properties/soap.client.props
```

AIX

Linux

UNIX

In the setupClient.sh file, the CLIENTSOAP variable is:

```
CLIENTSOAP=-Dcom.ibm.SOAP.ConfigURL=file:$WAS_HOME/properties/soap.client.props
```

In the setupClient.bat and the setupClient.sh files, complete the following steps:

1. Remove the leading slash (/) after file:.
2. Change sas to soap.

HP-UX

Java HotSpot Server VM warning: Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a message occurs in the native_stdout.log file when enabling security on the HP-UX11i platform

After you enable security on HP-UX 11i platforms, the following error in the native_stdout.log file occurs, along with a core dump and WebSphere Application Server does not start:

```
Java HotSpot(TM) Server VM warning:
Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a
```

To work around this error, apply the fixes recommended by Hewlett Packard for Java at the following URL: <http://www.hp.com/products1/unix/java/infocenter/patches.html>.

WebSphere Application Server Version 6 is not working correctly with Enterprise Workload Manager™ (EWLM)

To use WebSphere Application Server Version 6 with EWLM, you must manually update the WebSphere Application Server server.policy files. For example:

```
grant codeBase "file:/<EWLM_Install_Home>/classes/ARM/arm4.jar" {  
    permission java.security.AllPermission;  
};
```

Otherwise, you might encounter a Java 2 security exception for violating the Java 2 security permission.

For more information on configuring server.policy files, refer to the server.policy file permissions section in the *Developing and deploying applications* PDF book.

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

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

NMSV0610I: A NamingException is being thrown from a javax.naming.Context implementation

If you use CSiv2 inbound authentication, basic authentication is required, and Java clients running with com.ibm.CORBA.validateBasicAuth=true might fail with the following exception:

If you use CSiv2 inbound authentication, basic authentication is required, and Java™ clients running with com.ibm.CORBA.validateBasicAuth=true might fail with the following exception:

```
NMSV0610I: A NamingException is being thrown from a javax.naming.Context  
implementation. Details follow:
```

```
Context implementation: com.ibm.ws.naming.jndicos.CNContextImpl  
Context method: lookupExt  
Context name: TestaburgerNode01Cell/nodes/TestaburgerNode01/servers/server1  
Target name: SecurityServer  
Other data: ""  
Exception stack trace: javax.naming.NoPermissionException: NO_PERMISSION  
exception caught. Root exception is org.omg.CORBA.NO_PERMISSION:  
vmcid: 0x49421000 minor code: 92 completed: No  
...  
SECJ0395E: Could not locate the SecurityServer at host/port:9.42.72.27/9100  
to validate the userid and password entered. You may need to specify valid  
securityServerHost/Port in (WAS_INSTALL_ROOT)/properties/sas.client.props file.
```

To fix this problem, modify the com.ibm.CORBA.validateBasicAuth=false property in the client's sas.clients.props file, which is located in WAS_HOME/profiles/<profile-name>/properties, and then run the client.

When administrative security is enabled but application security is not enabled, the performance servlet displays authorization errors and cannot provide statistics

In WebSphere Application Server Version 6.1, when administrative security is enabled, the administration service is locked down. However, if application security is not enabled, an authentication challenge does not occur for incoming requests and, consequently, credentials do not exist for the performance servlet to access the administration service.

If administrative security is enabled, you also must enable application security for the performance servlet to process incoming requests.

"Name value is invalid" displays when migrating users and groups after the JACC provider for Tivoli is configured

When you use the migrateEAR utility to migrate the changes that were made to console users and groups after the JACC provider for Tivoli Access Manager is configured, the following configuration error displays in the systemOut.log file.

```
<specialSubjects> name value is invalid
```

The migrateEAR utility migrates the user and group data that is contained in the admin-authz.xml file. However, the migrateEAR utility does not convert the XML tags that are listed in the admin-authz.xml file if the pdwas-admin group is added to the administrator access control list (ACL) in Tivoli Access Manager prior to migration.

To resolve this error, enter the following command in padadmin to check whether the pdwas-admin group is in the administrator ACL before you migrate:

```
ac1 show
_WebAppServer_deployedResources_Roles_administrator_admin-Authz_ACL
```

The following result should display:

```
ACL Name:
_WebAppServer_deployedResources_Roles_administrator_admin-Authz_ACL
Description: Created by the Tivoli Access Manager
for Websphere Application Server Migration Tool.
Entries:
User sec_master TcmdbsvaBR1
Group pdwas-admin T[WebAppServer]i
```

If the pdwas-admin group is not listed, then enter the following command in padadmin to modify the ACL to add the pdwas-admin group:

```
ac1 modify
_WebAppServer_deployedResources_Roles_administrator_admin
-Authz_ACL set group pdwas-admin T [WebAppServer]i
```

A Sun JDK can not read a PKCS12 keystore created by the Application Server

A Sun JDK is not able to read a PKCS12 keystore created by the Application Server. The reason for this is that the PKCS12 implementation used by the IBM SDK and the Application Server is different than the implementation used by the Sun JDK. The difference causes problems when a Sun JDK is used to read the default trustore, trust.p12, or keystore, key.P12 created by the Application Server.

Because the truststore can not be read by the Sun JDK, you must first extract the certificates from the trustore using an IBM SDK. You can then import these certificates into a keystore that the Sun JDK can recognize correctly, such as a JKS keystore.

Security enablement followed by errors

Use this information if you are experiencing errors after security is enabled.

What kind of error are you seeing?

- Authentication error accessing a Web page
- Authorization error accessing a Web page
- "Authentication fails when code pages differ between the client and the server" on page 1304
- Error Message: CWSCJ0314E: Current Java 2 security policy reported a potential violation
- CWMSG0508E: The JMS Server security service was unable to authenticate user ID: error displayed in SystemOut.log when starting an application server
- Error Message: CWSCJ0237E: One or more vital LTPAServerObject configuration attributes are null or not available after enabling security and starting the application server

- An `AccessControlException` is reported in the `SystemOut.log`
- Error Message: `CWSCJ0336E: Authentication failed for user {0} because of the following exception {1}`
- “Error Message: `SECJ0352E: Could not get the users matching the pattern {0} because of the following exception {1}`” on page 1309
- “Generate keys error when using the Profile Management tool to create a new profile” on page 1309

For general tips on diagnosing and resolving security-related problems, see the topic [Troubleshooting the security component](#).

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/was/support/>

Authentication error accessing a Web page

Possible causes for authentication errors include:

- **Incorrect user name or passwords.** Check the user name and password and make sure that they are correct.
- **Security configuration error : User registry type is not set correctly.** Check the user registry property in administrative security settings in the administrative console. Verify that the user registry property is the intended user registry.
- **Internal program error.** If the client application is a Java standalone program, this program might not gather or send credential information correctly.

If the user registry configuration, user ID, and password appear correct, use the WebSphere Application Server trace to determine the cause of the problem. To enable security trace, use the `com.ibm.ws.security.*=all=enabled` trace specification.

Authorization error accessing a Web page

If a user who is supposed to have access to a resource does not, a configuration step is probably missing. For more information on configuring access to resources, review the chapter [Authorizing access to administrative roles](#) in the *Securing applications and their environment* PDF book.

Specifically:

- Check the required roles for the accessed Web resource.
- Check the authorization table to make sure that the user, or the groups to which the user belongs, is assigned to one of the required roles.
- View required roles for the Web resource in the deployment descriptor of the Web resource.
- View the authorization table for the application that contains the Web resource, using the administrative console.
- Test with a user who is granted the required roles, to see if the user can access the problem resources.
- If the user is required to have one or more of the required roles, use the administrative console to assign that user to required roles, stop, and restart the application.

If the user is granted required roles, but still fails to access the secured resources, enable security trace, using `com.ibm.ws.security.*=all=enabled` as the trace specification. Collect trace information for further resolution.

Authentication fails when code pages differ between the client and the server

When a client uses a code page that is different from the server, and non-US-ASCII characters are used for the user ID and password during basic authentication, the login does not succeed. The HTTP header

does not include the encoding method information that is necessary to translate the encoded data, so the server does not know how to decode the information correctly.

Use a login form that relies on POST parameters, which are in the HTML body text. The encoding for the text is sent by the browser and so is capable of being decoded properly.

Note: Web services customers are not able to use form login to resolve this problem. Users must ensure there is consistency in the code pages between the client and the server.

Error Message: CWSCJ0314E: Current Java 2 security policy reported a potential violation on server

If you find errors on your server similar to:

```
Error Message: CWSCJ0314E: Current Java 2 Security policy reported a potential violation of
Java 2 Security Permission. Please refer to Problem Determination Guide for further information.
{0}Permission/:{1}Code/:{2}{3}Stack Trace/:{4}Code Base Location/:{5}
```

The Java security manager checkPermission method has reported a SecurityException exception .

The reported exception might be critical to the secure system. Turn on security trace to determine the potential code that might have violated the security policy. Once the violating code is determined, verify if the attempted operation is permitted with respect to Java 2 Security, by examining all applicable Java 2 security policy files and the application code.

A more detailed report is enabled by either configuring RAS trace into debug mode, or specifying a Java property.

- Check the trace enabling section for instructions on how to configure Reliability Availability Serviceability (RAS) trace into debug mode, or
- Specify the following property in the **Application Servers > server_name > ProcessDefinition > Java Virtual Machine** panel from the administrative console in the **Generic JVM arguments** panel:
 - Add the **java.security.debug** run-time flag
 - Valid values:
 - access** Print all debug information including required permission, code, stack, and code base location.
 - stack** Print debug information including required permission, code, and stack.
 - failure** Print debug information including required permission, and code.

For a review of Java security policies, see the Java 2 Security documentation at <http://java.sun.com/j2se/1.3/docs/guide/security/index.html>.

Tip: If the application is running with a Java Mail application programming interface (API), this message might be benign. You can update the *installed Enterprise Application root/META-INF/was.policy* file to grant the following permissions to the application:

- permission java.io.FilePermission "\${user.home}\${/}.mailcap", "read";
- permission java.io.FilePermission "\${user.home}\${/}.mime.types", "read";
- permission java.io.FilePermission "\${java.home}\${/}lib\${/}mailcap", "read";
- permission java.io.FilePermission "\${java.home}\${/}lib\${/}mime.types", "read";

Error message: CWMSG0508E: The JMS Server security service was unable to authenticate user ID:" error displayed in SystemOut.log when starting an application server

This error can result from installing the Java Message Service (JMS) API sample and then enabling security. You can follow the instructions in the Configure and Run page of the corresponding JMS sample documentation to configure the sample to work with WebSphere Application Server security.

You can verify the installation of the message-driven bean sample by launching the installation program, selecting **Custom**, and browsing the components which are already installed in the Select the features you like to install panel. The JMS sample is shown as **Message-Driven Bean Sample**, under Embedded Messaging.

You can also verify this installation by using the administrative console to open the properties of the application server that contains the samples. Select **MDBSamples** and click uninstall.

Error message: CWSCJ0237E: One or more vital LTPAServerObject configuration attributes are null or not available after enabling security and starting the application server

This error message can result from selecting Lightweight Third Party Authentication (LTPA) as the authentication mechanism, but not generating the LTPA keys. The LTPA keys encrypt the LTPA token.

To resolve this problem:

1. Click **Security > Global security > Authentication > Authentication mechanisms and expiration > LTPA**
2. Enter a password, which can be anything.
3. Enter the same password in **Confirm Password**.
4. Click **Apply**.
5. Click **Generate Keys**.
6. Click **Save**.

The AccessControlException exception, is reported in the SystemOut.log

The problem is related to the Java 2 security feature of WebSphere Application Server, the API-level security framework that is implemented in WebSphere Application Server. An exception similar to the following example displays. The error message and number can vary.

```
CWSRV0020E: [Servlet Error]-[validator]: Failed to load servlet:  
java.security.AccessControlException: access denied  
(java.io.FilePermission  
app_server_root/systemApps/isclite.ear/isclite.war/WEB-INF/validation.xml read)
```

For an explanation of Java 2 security, how and why to enable or disable it, how it relates to policy files, and how to edit policy files, see the Java 2 security topic in the *Securing applications and their environment* PDF book. The topic explains that Java 2 security is not only used by this product, but developers can also implement it for their business applications. Administrators might need to involve developers, if this exception is created when a client tries to access a resource that is hosted by WebSphere Application Server.

Possible causes of these errors include:

- Syntax errors in a policy file.
- Syntax errors in permission specifications in the ra.xml file that is bundled in a .rar file. This case applies to resource adapters that support connector access to CICS® or other resources.
- An application is missing the specified permission in a policy file, or in permission specifications in ra.xml file bundled in a .rar file.
- The class path is not set correctly, preventing the permissions for the resource.xml file for Service Provider Programming Interface (SPI) from being correctly created.
- A library called by an application, or the application, is missing a doPrivileged block to support access to a resource.
- Permission is specified in the wrong policy file.

To resolve these problems:

- Check all of the related policy files to verify that the permission shown in the exception, for example java.io.FilePermission, is specified.

- Look for a related `ParserException` exception in the `SystemOut.log` file which reports the details of the syntax error.

For example:

```
CWSCJ0189E: Caught ParserException while creating template for application policy
```

```
profile_root/config/cells/cell_name/nodes/node_name/app.policy
```

Where:

- `cell_name` represents the name of your cell.
- `profile_name` represents the name of your profile.
- `node_name` represents the name of your node.

The exception is `com.ibm.ws.security.util.ParserException: line 18: expected ';', found 'grant'`

- Look for a message similar to: `CWSCJ0325W: The permission permission specified in the policy file is unresolved.`
- Check the call stack to determine which method does not have the permission. Identify the class path of this method. If it is hard to identify the method, enable the Java2 security Report.
 - Configuring RAS trace by specifying `com.ibm.ws.security.core.*=all=enabled`, or specifying a Java **property.java.security.debug** property. Valid values for the **java.security.debug** property are:
 - access** Print all debug information including: required permission, code, stack, and code base location.
 - stack** Print debug information including: required permission, code, and stack.
 - failure** Print debug information including: required permission and code.
 - The report shows:
 - Permission** The missing permission.
 - Code** Which method has the problem.
 - Stack Trace** Where the access violation occurred.
 - CodeBaseLocation** The detail of each stack frame.

Usually, permission and code are enough to identify the problem. The following example illustrates a report:

Permission:

```
app_server_root/logs/server1/SystemOut_02.08.20_11.19.53.log :
access denied (java.io.FilePermission
app_server_root/logs/server1/SystemOut_02.08.20_11.19.53.log delete)
```

Code:

```
com.ibm.ejs.ras.RasTestHelper$7 in
{file:app_server_root/installedApps/app1/JrasFVTApp.ear/RasLib.jar
}
```

Stack Trace:

```
java.security.AccessControlException: access denied (java.io.FilePermission
app_server_root/logs/server1/SystemOut_02.08.20_11.19.53.log delete
)
    at java.security.AccessControlContext.checkPermission
        (AccessControlContext.java(Compiled Code))
    at java.security.AccessController.checkPermission
        (AccessController.java(Compiled Code))
    at java.lang.SecurityManager.checkPermission
        (SecurityManager.java(Compiled Code))
    .
```

Code Base Location:

```
com.ibm.ws.security.core.SecurityManager :
file:/app_server_root/plugins/com.ibm.ws.runtime_6.1.0.jar
```

```
ClassLoader: com.ibm.ws.bootstrap.ExtClassLoader
Permissions granted to CodeSource
(file:/app_server_root/plugins/com.ibm.ws.runtime_6.1.0.jar <no certificates>
```

```
{
  (java.util.PropertyPermission java.vendor read);
  (java.util.PropertyPermission java.specification.version read);
  (java.util.PropertyPermission line.separator read);
  (java.util.PropertyPermission java.class.version read);
  (java.util.PropertyPermission java.specification.name read);
  (java.util.PropertyPermission java.vendor.url read);
  (java.util.PropertyPermission java.vm.version read);
  (java.util.PropertyPermission os.name read);
  (java.util.PropertyPermission os.arch read);
}
( This list continues.)
```

Where:

- *app1* represents the name of your application.
 - *app_server_root* represents the installation root directory for WebSphere Application Server Network Deployment.
 - *profile_root* represents the location and name of a particular profile in your system.
 - *profile1* or *profile_name* represents the name of your profile.
 - *server1* or *server_name* represents the name of your application server.
- If the method is SPI, check the resources.xml file to ensure that the class path is correct.
 - To confirm that all of the policy files are loaded correctly, or what permission each class path is granted, enable the trace with **com.ibm.ws.security.policy.*=all=enabled**. All loaded permissions are listed in the trace.log file. Search for the app.policy, was.policy and ra.xml files. To check the permission list for a class path, search for **Effective Policy for classpath**.
 - If there are any syntax errors in the policy file or the ra.xml file, correct them with the policy tool. Avoid editing the policy manually, because syntax errors can result. For additional information about using this tool, refer to the section Using PolicyTool to edit policy files in the *Developing and deploying applications* PDF book.
 - If a permission is listed as Unresolved, it does not take effect. Verify that the specified permission name is correct.
 - If the class path that is specified in the resource.xml file is not correct, correct it.
 - If a required permission does not exist in either the policy files or the ra.xml file, examine the application code to see if you need to add this permission. If so, add it to the proper policy file or the ra.xml file.
 - If the permission is not granted outside of the specific method that is accessing this resource, modify the code needs to use a doPrivileged block.
 - If this permission does exist in a policy file or a ra.xml file and the permission was loaded correctly, but the class path still does not have the permission in its list, the location of the permission might not be correct. Read the Java 2 security chapter in the *Securing applications and their environment* PDF book carefully to determine in which policy file or ra.xml file to specify that permission.

Tip: If the application is running with the Java Mail API, you can update the *installed Enterprise Application root/META-INF/was.policy* file to grant the following permissions to the application:

- permission java.io.FilePermission "\${user.home}/\${}.mailcap", "read";
- permission java.io.FilePermission "\${user.home}/\${}.mime.types", "read";
- permission java.io.FilePermission "\${java.home}/\${}lib\${}/mailcap", "read";
- permission java.io.FilePermission "\${java.home}/\${}lib\${}/mime.types", "read";

Error Message: CWSCJ0336E: Authentication failed for user {0} because of the following exception {1}

This error message results if the user ID that is indicated is not found in the Lightweight Directory Access Protocol (LDAP) user registry. To resolve this problem:

1. Verify that your user ID and password are correct.
2. Verify that the user ID exists in the registry.

3. Verify that the base distinguished name (DN) is correct.
4. Verify that the user filter is correct.
5. Verify that the bind DN and the password for the bind DN are correct. If the bind DN and password are not specified, add the missing information and retry.
6. Verify that the host name and LDAP type are correct.

Consult with the administrator of the user registry if the problem persists.

Error Message: SECJ0352E: Could not get the users matching the pattern {0} because of the following exception {1}

This authentication failure message displays when an external user account repository is corrupted or unavailable, and WebSphere Application Server is unable to authenticate the user name in the repository. Generally, authentication error messages are followed by additional information that indicates the nature or root cause of the problem, such as:

Make sure the users matching the pattern exist in the registry. Contact your service representative if the problem persists.

This additional information might not provide a clear user action if the user account repository is corrupted or the user loses connectivity between WebSphere Application Server and an external user account repository. The external user account repository, which is referred to as a repository in this document, might be a Lightweight Directory Access Protocol (LDAP) product.

To resolve this problem, you might need to re-install the repository and verify that it installs successfully by testing the connection.

Note: Proceed with the following steps only if you have ensured that all WebSphere Application Server-related configuration settings are accurate.

Complete the following steps to resolve the issue:

1. Restart both the repository and WebSphere Application Server.
2. Test the connection to the repository. If the connection attempt still fails, it might be necessary to re-install the repository.
3. If diagnostics are provided with the repository, run them to avoid having to re-install the repository.

Note: If the previous steps do not fix the problem, you might need to re-install the repository. Before proceeding, generate a complete list of all the configured users and groups; you will need to re-populate these fields after the re-installation.

4. If necessary, re-install the corrupted repository.
5. Populate the users and groups from your list into the newly installed repository.
6. Restart both the repository and WebSphere Application server.
7. In the administrative console, navigate to **Security > Global security**, and select the appropriate user account repository. For example, select **Standalone LDAP registry** if you are using a standalone Lightweight Directory Access Protocol repository.
8. Click **Test connection** to ensure that WebSphere Application Server can connect to the repository.

Generate keys error when using the Profile Management tool to create a new profile

When you create a new profile using either the Profile Management tool or the command-line manageprofiles utility, an error message displays that indicates either partial success or failure. The error message, which is located in the *install_dir/logs/manageprofiles/profile_name_create.log* file, might point to an error in either the generateKeysforSingleProfile task or the generateKeysForCellProfile task.

The Profile Creation tool and the manageprofiles utility invoke several tasks. The generateKeysforSingleProfile task is invoked when you create a stand-alone application server or a

deployment manager profile. The generateKeysForCellProfile task is invoked when you create a cell profile. Both of these tasks are the first tasks to invoke the wsadmin commands. Although the log indicates an error in one of these tasks, the error might actually result from a wsadmin command failure and not an error in the security tasks.

To determine the actual cause of the problem, review the information that is provided in the following log files:

- *install_dir/logs/manageprofiles/profile_name_create.log* file indicates the error code of the failure
- *install_dir/logs/manageprofiles/profile_name/keyGeneration.log* file
- *install_dir/logs/manageprofiles/profile_name/wsadminListener.log* file

Access problems after enabling security

Use this information if you are experiencing access problems after enabling security.

What kind of error are you seeing?

- “I cannot access all or part of the administrative console or use the wsadmin tool after enabling security”
- “I cannot access a Web page after enabling security” on page 1311
- “Authentication error accessing a Web page” on page 1304
- “Authorization error accessing a Web page” on page 1304
- “The client cannot access an enterprise bean after enabling security” on page 1311
- “Client program never gets prompted when accessing secured enterprise bean” on page 1313
- “Cannot stop an application server, node manager, or node after enabling security” on page 1313
- “The AccessControlException exception, is reported in the SystemOut.log” on page 1306
- “After enabling single sign-on, I cannot logon to the administrative console” on page 1314
- “The following exception displays in the SystemOut.log file after I start the server and enable security: “SECJ0306E: No received or invocation credential exists on the thread.” ” on page 1314
- “A Name NotFoundException error occurs when initially connecting to the federated repositories.” on page 1314

For general tips on diagnosing and resolving security-related problems, see the topic “Security components troubleshooting tips” on page 1289.

If you do not see a problem that resembles yours, or if the information provided does not solve your problem, see Troubleshooting help from IBM.

I cannot access all or part of the administrative console or use the wsadmin tool after enabling security

- If you cannot access the administrative console, or view and update certain objects, look in the SystemOut log of the application server which hosts the administrative console page for a related error message.
- You might not have authorized your ID for administrative tasks. This problem is indicated by errors such as:
 - [8/2/02 10:36:49:722 CDT] 4365c0d9 RoleBasedAuth A CWSCJ0305A: Role based authorization check failed for security name MyServer/myUserId, accessId MyServer/S-1-5-21-882015564-4266526380-2569651501-1005 while invoking method getProcesType on resource Server and module Server.
 - Exception message: “CWWMN0022E: Access denied for the getProcesType operation on Server MBean”
 - When running the command: wsadmin -username j2ee -password j2ee: CWWAX7246E: Cannot establish “SOAP” connection to host “BIRKT20” because of an authentication failure. Ensure that user and password are correct on the command line or in a properties file.

To grant an ID administrative authority, from the administrative console, click **System Administration > Console Users** and validate that the ID is a member. If the ID is not a member, add the ID with at least monitor access privileges, for read-only access.

- Verify that the `enable_trusted_application` flag is set to `true`. To check the `enable_trusted_application` flag value using the administrative console, click **Security > Global security**. Under Additional properties, click **Custom properties > EnableTrustedApplications**.

I cannot access a Web page after enabling security

When secured resources are not accessible, probable causes include:

- Authentication errors - WebSphere Application Server security cannot identify the ID of the person or process. Symptoms of authentication errors include:
 - On a Netscape browser:
 - Authorization failed. Retry? message is displayed after an attempt to log in.
 - Accepts any number of attempts to retry login and displays Error 401 message when Cancel is clicked to stop retry.
 - A typical browser message displays: Error 401: Basic realm='Default Realm'.
 - On an Internet Explorer browser:
 - Login prompt displays again after an attempt to log in.
 - Allows three attempts to retry login.
 - Displays Error 401 message after three unsuccessful retries.
- Authorization errors - The security function has identified the requesting person or process as not authorized to access the secured resource. Symptoms of authorization errors include:
 - Netscape browser: "Error 403: AuthorizationFailed" message is displayed.
 - Internet Explorer:
 - "You are not authorized to view this page" message is displayed.
 - "HTTP 403 Forbidden" error is also displayed.
- SSL errors - WebSphere Application Server security uses Secure Sockets Layer (SSL) technology internally to secure and encrypt its own communication, and incorrect configuration of the internal SSL settings can cause problems. Also you might have enabled SSL encryption for your own Web application or enterprise bean client traffic which, if configured incorrectly, can cause problems regardless of whether WebSphere Application Server security is enabled.
 - SSL-related problems are often indicated by error messages that contain a statement such as:


```
ERROR: Could not get the initial context or unable to look up the starting context.Exiting. followed by javax.net.ssl.SSLHandshakeException
```

The client cannot access an enterprise bean after enabling security

If the client access to an enterprise bean fails after security is enabled:

- Review the steps for securing and granting access to resources.
- Browse the server JVM logs for errors relating to enterprise bean access and security. Look up any errors in the message table.

Errors similar to Authorization failed for /UNAUTHENTICATED while invoking *resource* `securityName:/UNAUTHENTICATED;accessId:UNAUTHENTICATED` not granted any of the required roles *roles* indicate that:

- An unprotected servlet or JavaServer Pages (JSP) file accessed a protected enterprise bean. When an unprotected servlet is accessed, the user is not prompted to log in and the servlet runs as UNAUTHENTICATED. When the servlet makes a call to an enterprise bean that is protected, the servlet fails.

To resolve this problem, secure the servlet that is accessing the protected enterprise bean. Make sure that the `runAs` property for the servlet is set to an ID that can access the enterprise bean.

- An unauthenticated Java client program is accessing an enterprise bean resource that is protected. This situation can happen if the file that is read by the `sas.client.props` properties file that is used by the client program does not have the `securityEnabled` flag set to `true`.

To resolve this problem, make sure that the `sas.client.props` file on the client side has its `securityEnabled` flag set to `true`.

Errors similar to Authorization failed for *valid_user* while invoking *resource* securityName:/username;accessId:xxxxxx not granted any of the required roles *roles* indicate that a client attempted to access a secured enterprise bean resource, and the supplied user ID is not assigned the required roles for that enterprise bean.

- Check that the required roles for the enterprise bean resource are accessed. View the required roles for the enterprise bean resource in the deployment descriptor of the Web resource.
- Check the authorization table and make sure that the user or the group that the user belongs to is assigned one of the required roles. You can view the authorization table for the application that contains the enterprise bean resource using the administrative console.

If org.omg.CORBA.NO_PERMISSION exceptions occur when programmatically logging on to access a secured enterprise bean, an authentication exception has occurred on the server. Typically the CORBA exception is triggered by an underlying com.ibm.WebSphereSecurity.AuthenticationFailedException. To determine the actual cause of the authentication exception, examine the full trace stack:

1. Begin by viewing the text following WSSecurityContext.acceptSecContext(), reason: in the exception. Typically, this text describes the failure without further analysis.
2. If this action does not describe the problem, look up the Common Object Request Broker Architecture (CORBA) minor code. The codes are listed in the article titled Troubleshooting the security components reference.

For example, the following exception indicates a CORBA minor code of 49424300. The explanation of this error in the CORBA minor code table reads:

```
authentication failed error
```

In this case the user ID or password supplied by the client program is probably not valid:

```
org.omg.CORBA.NO_PERMISSION: Caught WSSecurityContextException in
WSSecurityContext.acceptSecContext(), reason: Major Code[0] Minor Code[0]
Message[ Exception caught invoking authenticateBasicAuthData from SecurityServer
for user jdoe. Reason: com.ibm.WebSphereSecurity.AuthenticationFailedException]
minor code: 49424300 completed:
No at com.ibm.ISecurityLocalObjectBaseL13Impl.PrincipalAuthFailReason.map_auth_fail_to_minor_code
(PrincipalAuthFailReason.java:83)
```

A CORBA INITIALIZE exception with CWSA1477W: SECURITY CLIENT/SERVER CONFIGURATION MISMATCH error embedded, is received by client program from the server.

This error indicates that the security configuration for the server differs from the client in some fundamental way. The full exception message lists the specific mismatches. For example, the following exception lists three errors:

```
Exception received: org.omg.CORBA.INITIALIZE:
CWSA1477W: SECURITY CLIENT/SERVER CONFIG MISMATCH:
The client security configuration (sas.client.props or outbound settings in
administrative console) does not support the server security configuration for
the following reasons:
ERROR 1: CWSA0607E: The client requires SSL Confidentiality but the server does not
support it.
ERROR 2: CWSA0610E: The server requires SSL Integrity but the client does not
support it.
ERROR 3: CWSA0612E: The client requires client (e.g., userid/password or token),
but the server does not support it.
minor code: 0
completed: No at
com.ibm.ISecurityLocalObjectBaseL13Impl.SecurityConnectionInterceptor.getConnectionKey
(SecurityConnectionInterceptor.java:1770)
```

In general, resolving the problem requires a change to the security configuration of either the client or the server. To determine which configuration setting is involved, look at the text following the CWSA error message. For more detailed explanations and instructions, look in the message reference, by selecting the **Reference** view of the information center navigation and expanding **Messages** in the navigation tree.

In these particular cases:

- In ERROR 1, the client is requiring SSL confidentiality but the server does not support SSL confidentiality. Resolve this mismatch in one of two ways. Either update the server to support SSL confidentiality or update the client so that it no longer requires it.
- In ERROR 2, the server requires SSL integrity but the client does not support SSL integrity. Resolve this mismatch in one of two ways. Either update the server to support SSL integrity or update the client so that it no longer requires it.
- In ERROR 3, the client requires client authentication through a user id and password, but the server does not support this type of client authentication. Either the client or the server needs to change the configuration. To change the client configuration, modify the `SAS.CLIENT.PROPS` file for a pure client or change the outbound configuration for the server in the Security administrative console. To change the configuration for the target server, modify the inbound configuration in the Security administrative console.

Similarly, an exception like `org.omg.CORBA.INITIALIZE: JSAS0477W: SECURITY CLIENT/SERVER CONFIG MISMATCH`: appearing on the server trying to service a client request indicates a security configuration mismatch between client and server. The steps for resolving the problem are the same as for the `JSAS1477W` exceptions previously described.

Client program never gets prompted when accessing secured enterprise bean

Even though it seems that security is enabled and an enterprise bean is secured, occasions can occur when the client runs the remote method without prompting. If the remote method is protected, an authorization failure results. Otherwise, run the method as an unauthenticated user.

Possible reasons for this problem include:

- The server with which you are communicating might not have security enabled. Check with the WebSphere Application Server administrator to ensure that the server security is enabled. Access the security settings from within the **Security** section of the administrative console.
- The client does not have security enabled in the `sas.client.props` file. Edit the `sas.client.props` file to ensure the property `com.ibm.CORBA.securityEnabled` is set to `true`.
- The client does not have a `ConfigURL` specified. Verify that the property `com.ibm.CORBA.ConfigURL` is specified on the command line of the Java client, using the `-D` parameter.
- The specified `ConfigURL` does not have a valid URL syntax, or the `sas.client.props` that is pointed to cannot be found. Verify that the `com.ibm.CORBA.ConfigURL` property is valid. Check the Java documentation for a description of URL formatting rules. Also, validate that the file exists at the specified path.

Windows An example of a valid property is `C:/WebSphere/AppServer/properties/sas.client.props`.

- The client configuration does not support message layer client authentication (user ID and password). Verify that the `sas.client.props` file has one of the following properties set to `true`:
 - `com.ibm.CSI.performClientAuthenticationSupported=true`
 - `com.ibm.CSI.performClientAuthenticationRequired=true`
- The server configuration does not support message layer client authentication, which consists of a user ID and password. Check with the WebSphere Application Server administrator to verify that user ID and password authentication is specified for the inbound configuration of the server within the System Administration section of the administrative console administration tool.

Cannot stop an application server, node manager, or node after enabling security

If you use command-line utilities to stop WebSphere Application Server processes, apply additional parameters after enabling security to provide authentication and authorization information.

Use the `./stopServer -help` command to display the parameters to use.

Use the following command options after enabling security:

- `./stopServer serverName -username name -password password`
- `./stopNode -username name -password password`
- `./stopManager -username name -password password`

If you use the Windows service panel or the `net stop` command to stop the WebSphere Application Server processes and the service could not be stopped, update the existing Application Server service using additional stop arguments. You might need to end the server process from the Task Manager before updating the service. Use the `-stopArgs` and the `-encodeParams` parameters to update the service as described in the "Updating an existing Application Server service" example in the WASService command chapter of the *Administering applications and their environment* PDF book..

After enabling single sign-on, I cannot logon to the administrative console

This problem occurs when single sign-on (SSO) is enabled, and you attempt to access the administrative console using the short name of the server, for example `http://myserver:port_number/ibm/console`. The server accepts your user ID and password, but returns you to the logon page instead of the administrative console.

To correct this problem, use the fully qualified host name of the server, for example `http://myserver.mynetwork.mycompany.com:9060/ibm/console`.

The following exception displays in the SystemOut.log file after I start the server and enable security: "SECJ0306E: No received or invocation credential exists on the thread."

The following message displays when one or more nodes within the cell was not synchronized during configuration:

```
SECJ0306E: No received or invocation credential exists on the thread. The Role based
authorization check will not have an accessId of the caller to check. The parameters
are: access check method getServerConfig on resource FileTransferServer and module
FileTransferServer. The stack trace is java.lang.Exception: Invocation and received
credentials are both null.
```

Make sure that each of the nodes are synchronized and then restart the deployment manager.

A NameNotFoundException error occurs when initially connecting to the federated repositories.

When the server attempts an indirect lookup on the `java:comp/env/ds/wimDS` name and makes its initial EJB connection to the federated repositories, the following error message displays in the SystemOut.log file:

```
NMSV0612W: A NameNotFound Exception
```

The `NameNotFoundException` error is caused by the reference binding definition for the `jdbc/wimDS` Java Naming and Directory interface (JNDI) name in the `ibm-ejb-jar-bnd.xmi` file. You can ignore this warning message. The message does not display when the `wimDS` database repository is configured.

Secure Sockets Layer errors

You might encounter various problems after configuring or enabling Secure Sockets Layer (SSL). You may not be able to stop the deployment manager after configuring the SSL. You may not be able to access resource using HTTPS. The client and the server may not be able to negotiate the proper level of security. The problems mentioned here are only a few of the possibilities. Solving these problems is imperative to the successful operation of WebSphere Application Server.

What type of problem are you having?

- “Stopping the deployment manager after configuring Secure Sockets Layer”
- “Accessing resources using HTTPS”
- “javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: handshake failure” on page 1316
- “javax.net.ssl.SSLHandshakeException: unknown certificate” on page 1316
- “javax.net.ssl.SSLHandshakeException: bad certificate” on page 1317
- “org.omg.CORBA.INTERNAL: EntryNotFoundException or NTRegistryImp E CWSCJ0070E: No privilege id configured for: error when programmatically creating a credential” on page 1317
- ““Catalog” tablet is blank (no item displayed) in GUI application client” on page 1318
- “Modifying SSL Configurations after migration using -scriptCompatibility true” on page 1318

Stopping the deployment manager after configuring Secure Sockets Layer

After configuring the Secure Sockets Layer repertoires, if you stop the deployment manager without also stopping the node agents, you might receive the following error message when you restart the deployment manager:

```
CWWMU0509I: The server "nodeagent" cannot be reached. It appears to be stopped.
CWWMU0211I: Error details may be seen in the file:
/opt/WebSphere/AppServer/logs/nodeagent/stopServer.log
```

The error occurs because the deployment manager did not propagate the new SSL certificate to the node agents. The node agents are using an older certificate file than the deployment manager and the certificate files are incompatible. To work around this problem, you must manually stop the node agent and deployment manager processes.

Windows To end the processes, use the Task Manager.

UNIX Run the command to end the process

You need to consider certain items when identifying the specific process to stop. For each process that is stopped, WebSphere Application Server stores the process ID in a pid file and you need to find these *.pid files. For example, the server1.pid for a stand-alone install action might be found at: *install_root/logs/server1.pid*

Accessing resources using HTTPS

If you are unable to access resources using a Secure Sockets Layer (SSL) URL (beginning with https:), or encounter error messages that indicate SSL problems, verify that your HTTP server is configured correctly for SSL. Browse the welcome page of the HTTP server using SSL by entering the URL: **https://host_name.**

If the page works with HTTP, but not HTTPS, the problem is with the HTTP server.

- Refer to the documentation for your HTTP server for instructions on correctly enabling SSL. If you are using the IBM HTTP Server or Apache, go to: <http://www.ibm.com/software/webservers/httpservers/library.html>. Click **Frequently Asked Questions > SSL**.
- If you use the IBM Key Management (IKeyman) tool to create certificates and keys, remember to stash the password to a file when creating the Key Database (KDB) file with the IBM Key Management Tool.
 1. Go to the directory where the KDB file is created, and see if an .sth file exists.
 2. If not, open the KDB file with the IBM Key Management Tool, and click **Key Database File > Stash Password**. The following message is displayed: The password has been encrypted and saved in the file.

If the HTTP server handles SSL-encrypted requests successfully, or is not involved (for example, traffic flows from a Java client application directly to an enterprise bean that is hosted by WebSphere Application Server, or the problem displays only after enabling WebSphere Application Server security), what kind of error are you seeing?

- javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: handshake failure
- javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: unknown certificate
- javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: bad certificate

You get this error message org.omg.CORBA.INTERNAL: EntryNotFoundException or NRegistryImp E CWSCJ0070E: No privilege id configured for: when programmatically creating a credential

For general tips on diagnosing and resolving security-related problems, see “Security components troubleshooting tips” on page 1289

If you do not see a problem that resembles yours, or if the information provided does not solve your problem, see Troubleshooting help from IBM

javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: handshake failure

If you see a Java exception stack similar to the following example:

```
[Root exception is org.omg.CORBA.TRANSIENT: CAUGHT_EXCEPTION_WHILE_CONFIGURING_
SSL_CLIENT_SOCKET: CWWJE0080E: javax.net.ssl.SSLHandshakeException - The client
and server could not negotiate the desired level of security. Reason: handshake
failure:host=MYSERVER,port=1079 minor code: 4942F303 completed: No] at
com.ibm.CORBA.transport.TransportConnectionBase.connect
(TransportConnectionBase.java:NNN)
```

Some possible causes are:

- Not having common ciphers between the client and server.
- Not specifying the correct protocol.

To correct these problems:

1. Review the SSL settings. In the administrative console, click **Security > SSL certificate and key management**. Under Configuration settings, click **Manage endpoint security configurations > endpoint_configuration_name**. Under Related items, click **SSL configurations > SSL_configuration_name**. You can also browse the file manually by viewing the *install_root/properties/sas.client.props* file.
2. Check the property that is specified by the *com.ibm.ssl.protocol* file to determine which protocol is specified.
3. Check the cipher types that are specified by the *com.ibm.ssl.enabledCipherSuites* interface. You might want to add more cipher types to the list. To see which cipher suites are currently enabled, click **Quality of protection settings (QoP)**, and look for the **Cipher Suites** property.
4. Correct the protocol or cipher problem by using a different client or server protocol and cipher selection. Typical protocols are SSL or SSLv3.
5. Make the cipher selection 40-bit instead of 128-bit. For Common Secure Interoperability Version 2 (CSlv2), set both of the following properties to false in the *sas.client.props* file, or set security level=medium in the administrative console settings:
 - *com.ibm.CSI.performMessageConfidentialityRequired*=false
 - *com.ibm.CSI.performMessageConfidentialitySupported*=false

javax.net.ssl.SSLHandshakeException: unknown certificate

If you see a Java exception stack similar to the following example, it might be caused by not having the personal certificate for the server in the client truststore file:

```
ERROR: Could not get the initial context or unable to look up the starting context.
Exiting. Exception received: javax.naming.ServiceUnavailableException: A
communication failure occurred while attempting to obtain an initial context using
```

the provider url: "corbaloc:iiop:localhost:2809". Make sure that the host and port information is correct and that the server identified by the provider url is a running name server. If no port number is specified, the default port number 2809 is used. Other possible causes include the network environment or workstation network configuration. [Root exception is org.omg.CORBA.TRANSIENT: CAUGHT_EXCEPTION_WHILE_CONFIGURING_SSL_CLIENT_SOCKET: CWWJE0080E: javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: unknown certificate:host=MYSERVER,port=1940 minor code: 4942F303 completed: No]

To correct this problem:

1. Check the client truststore file to determine if the signer certificate from the server personal certificate is there. For a self-signed server personal certificate, the signer certificate is the public key of the personal certificate. For a certificate authority (CA)-signed server personal certificate, the signer certificate is the root CA certificate of the CA that signed the personal certificate.
2. Add the server signer certificate to the client truststore file.

javax.net.ssl.SSLHandshakeException: bad certificate

A Java exception stack error might display if the following situations occur:

- A personal certificate exists in the client keystore that is used for SSL mutual authentication.
- The signer certificate is not extracted into the server truststore file, and thus the server cannot trust the certificate whenever the SSL handshake is made.

The following message is an example of the Java exception stack error:

```
ERROR: Could not get the initial context or unable to look
up the starting context. Exiting.
Exception received: javax.naming.ServiceUnavailableException:
A communication failure occurred while attempting to obtain an
initial context using the provider url: "corbaloc:iiop:localhost:2809".
Make sure that the host and port information is correct and that the
server identified by the provider url is a running name
server. If no port number is specified, the default port number 2809
is used. Other possible causes include the network environment or
workstation network configuration.
[Root exception is org.omg.CORBA.TRANSIENT: CAUGHT_EXCEPTION_WHILE_CONFIGURING_SSL_
CLIENT_SOCKET: CWWJE0080E: javax.net.ssl.SSLHandshakeException - The client and
server could not negotiate the desired level of security. Reason:
bad certificate: host=MYSERVER,port=1940 minor code: 4942F303 completed: No]
```

To verify this problem, check the server truststore file to determine if the signer certificate from the client personal certificate is there. For a self-signed client personal certificate, the signer certificate is the public key of the personal certificate. For a certificate authority-signed client personal certificate, the signer certificate is the root CA certificate of the CA that signed the personal certificate.

To correct this problem, add the client signer certificate to the server truststore file.

org.omg.CORBA.INTERNAL: EntryNotFoundException or NTRegistryImp E CWSJC0070E: No privilege id configured for: error when programmatically creating a credential

If you encounter the following exception in a client application attempting to request a credential from a WebSphere Application Server using SSL mutual authentication:

```
ERROR: Could not get the initial context or unable to look up the starting context.
Exiting. Exception received: org.omg.CORBA.INTERNAL: Trace from server: 1198777258
at host MYHOST on port 0 >>org.omg.CORBA.INTERNAL: EntryNotFoundException minor
code: 494210B0 completed:
No at com.ibm.ISecurityLocalObjectBaseL13Impl.PrincipalAuthFailReason.
map_auth_fail_to_minor_code(PrincipalAuthFailReason.java:99)
```

or a simultaneous error from the WebSphere Application Server that resembles:

```
[7/31/02 15:38:48:452 CDT] 27318f5 NRegistryImp E CWSCJ0070E: No privilege id
configured for: testuser
```

The cause might be that the user ID sent by the client to the server is not in the user registry for that server.

To confirm this problem, check that an entry exists for the personal certificate that is sent to the server. Depending on the user registry mechanism, look at the native operating system user ID or Lightweight Directory Access Protocol (LDAP) server entries.

To correct this problem, add the user ID to the user registry entry (for example, operating system, LDAP directory, or other custom registry) for the personal certificate identity.

"Catalog" tablet is blank (no item displayed) in GUI application client

This error message occurs when you install an ActiveX client sample application that uses the PlantsByWebSphere Active X to EJB Bridge.

The cause is that the server certificate is not in the client trustore that is specified in the client.ssl.props file. Although the "com.ibm.ssl.enableSignerExchangePrompt" signer property might be set to true, the auto-exchange prompt only supports a command-line prompt. If the sample application relies on a graphical user interface and does not provide access to a command prompt, for example using standard in and standard out, the auto-exchange prompt does not function.

Note: The applet client under the Client Technology Samples does not have access to the command prompt and it cannot see the auto-exchange prompt. Thus, the applet client cannot rely on the auto-exchange prompt feature.

To correct this problem, retrieve the certificate manually using the retrieveSigners utility.

Modifying SSL Configurations after migration using -scriptCompatibility true

After migrating using scriptCompatibility true, all attributes of the SSL configurations cannot be edited through the administrative console. In particular, the hardware cryptography settings cannot be displayed or edited.

By using the scriptCompatibility true flag, the SSL configurations are not migrated to the new format for support in the 6.1 release. New capabilities were added that are not supported when the configurations are not migrated to the latest format.

To solve this problem, create new SSL configuration definitions to replace those in the pre 6.1 format or continue to edit those configurations using scripting.

Errors configuring Secure Sockets Layer encrypted access

You might have errors returned when you are trying to configure Secure Sockets Layer (SSL) for encrypted access. This article describes some of the common errors you might encounter and makes suggestions on how to fix the problems.

What kind of error are you seeing?

- "'The Java Cryptographic Extension (JCE) files were not found.'" error when launching iKeyman" on page 1319
- "'Unable to verify MAC.'" error when the wrong keystore password is used" on page 1319
- "'SSL handshake failure" error when no trusted certificate is found" on page 1319
- "The certificate alias cannot be found in the keystore" on page 1320

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.

"The Java Cryptographic Extension (JCE) files were not found." error when launching iKeyman

You might receive the following error when you attempt to start the iKeyman tool:

```
"The Java Cryptographic Extension (JCE) files were not found.  
Please check that the JCE files have been installed in the correct directory."
```

When you click **OK**, the iKeyman tool closes. To resolve this problem:

- Set the `JAVA_HOME` parameter so that it points to the Java Developer Kit that is shipped with WebSphere Application Server.

UNIX For example, the command is similar to: `export JAVA_HOME=/opt/WebSphere/AppServer/java`

Windows If WebSphere Application Server is installed on your c: drive, the command would be: `set JAVA_HOME=c:\WebSphere\AppServer\java`

- Rename the file `install_dir/java/jre/lib/ext/gskikm.jar` to `gskikm.jar.org`.

The file is located in the `install_dir/java/jre/lib/ext/` directory.

By default, the file is located in the following directory: `app_server_rootedition_name/java/ext`.

"Unable to verify MAC." error when the wrong keystore password is used

You might receive the following error when the keystore password is not being used correctly.

```
CWPKI0033E: The keystore located at "C:/WebSphere/AppServer/profiles/AppSrv01/etc/trust.p12"  
failed to load due to the following error: Unable to verify MAC.
```

Change the **Password** field that references this keystore by using the correct password. The default password is `WebAS`. Never use this password in a production environment.

"SSL handshake failure" error when no trusted certificate is found

You might receive the following error when you attempt to add the signer to the local truststore:

```
CWPKI0022E: SSL HANDSHAKE FAILURE: A signer with SubjectDN "CN=BIRKT40.austin.ibm.com,  
O=IBM, C=US" was sent from target host:port "9.65.49.131:9428".
```

The signer might need to be added to the local truststore `C:/WASX_c0602.31/AppServer/profiles/Dmgr09/etc/trust.p12` that is located in the SSL configuration alias `DefaultSSLSettings`. The truststore is loaded from the SSL configuration file.

The extended error message from the SSL handshake exception is:

```
"No trusted certificate found."
```

This error indicates that the signer certificate from the specified target host and port has not been located in the specified truststore, the SSL settings, and the SSL configuration file. If this occurs in a client process, there are several things that you can do:

- Enable the signer exchange prompt.
- Run the **retrieveSigners** script. For more information about the location of the signer certificate see the `retrieveSigners` command topic in the *Securing applications and their environment PDF* book.
- Manually export the signers from the server and import them to the client.

If this issue occurs in a server process, then complete one of the following procedures:

- In the administrative console, find the target endpoint (host name and port) and determine the certificate that is used by the SSL configuration associated with it. Extract the SSL certificate to a file, and import it into the sending server truststore that is referenced in the error message.
- In the administrative console, find the sending server truststore. Go to signer certificates, add from Port, and connect directly to the target host and port, which are indicated in the message, to retrieve the signer directly into the truststore.
- Manually extract the signer from the target server and host keystore by using the iKeyman utility, and import the signer into the truststore of the server sending the certificate.

Note: As default, Websphere Application Server uses the key.p12 and trust.p12 files for any communication between Websphere Application Servers (for example between nodeagent and appserver or vice versa). If WebSphere Application Server is looking for a certificate in some file other than these, then it is possible that your application establishes the secure socket layer (SSL) configuration by using system properties, established with the System.setProperty() method. That is, SSL configurations are managed for each of your processes, and you have had to maintain individual settings for each SSL configuration in the topology.

Prior to WebSphere Application Server Version 6.1, the WebSphere Application Server management processes allowed the individually-managed SSL configurations which were set by system properties. Your pre-Version 6.1 system properties settings were processed successfully.

With WebSphere Application Server Version 6.1, central management of Secure Sockets Layer (SSL) configuration occurs. Applications that use SSL connections based on values set for system properties instead of using the centrally managed default dynamic SSL configuration can experience handshake failures. Nodeagent to appserver communications is being governed by the default dynamic SSL configuration in WebSphere Application Server Version 6.1 and not through the system properties you set. You may need to adjust your application to use the centrally managed SSL configuration of WebSphere Application Server Version 6.1.

The certificate alias cannot be found in the keystore

You might receive the following error when the certificate alias is not found in the referenced keystore:

```
CWPKI0023E: The certificate alias "default" specified by the property
com.ibm.ssl.keyStoreClientAlias is not found in KeyStore
"c:/WebSphere/AppServer/profiles/Dmgr01/config/cells/myCell/key.p12".
```

This error indicates that the certificate alias that was specified cannot be found in the referenced keystore. Either change the certificate alias or make sure that alias exists in the specified keystore.

Single sign-on configuration troubleshooting tips

Several common problems can occur when you configure single sign-on (SSO) between a WebSphere Application Server and a Domino server. Some such problems are: Failure to save the Domino Web SSO configuration, authentication failures when accessing a protected resource, and SSO failures when accessing a protected resource. You can take some actions to correct these error situations and restore the SSO.

- Failure to save the Domino Web SSO configuration document

The client must find Domino server documents for the participating SSO Domino servers. The Web SSO configuration document is encrypted for the servers that you specify. The home server that is indicated by the client location record must point to a server in the Domino domain where the participating servers reside. This pointer ensures that lookups can find the public keys of the servers.

If you receive a message stating that one or more of the participating Domino servers cannot be found, then those servers cannot decrypt the Web SSO configuration document or perform SSO.

When the Web SSO configuration document is saved, the status bar indicates how many public keys are used to encrypt the document by finding the listed servers, authors, and administrators in the document.

- Failure of the Domino server console to load the Web SSO configuration document at Domino HTTP server startup

During configuration of SSO, the server document is configured for Multi-Server in the **Session Authentication** field. The Domino HTTP server tries to find and load a Web SSO configuration document during startup. The Domino server console reports the following information if a valid document is found and decrypted: HTTP: Successfully loaded Web SSO Configuration.

If a server cannot load the Web SSO configuration document, SSO does not work. In this case, a server reports the following message: HTTP: Error Loading Web SSO configuration. Reverting to single-server session authentication.

Verify that only one Web SSO configuration document is in the Web configurations view of the Domino directory and in the \$WebSSOConfigs hidden view. You cannot create more than one document, but you can insert additional documents during replication.

If you can verify only one Web SSO configuration document, consider another condition. When the public key of the server document does not match the public key in the ID file, this same error message can display. In this case, attempts to decrypt the Web SSO configuration document fail and the error message is generated.

This situation can occur when the ID file is created multiple times, but the Server document is not updated correctly. Usually, an error message is displayed on the Domino server console stating that the public key does not match the server ID. If this situation occurs, SSO does not work because the document is encrypted with a public key for which the server does not possess the corresponding private key.

To correct a key-mismatch problem:

1. Copy the public key from the server ID file and paste it into the Server document.
2. Create the Web SSO configuration document again.

- Authentication fails when accessing a protected resource.

If a Web user is repeatedly prompted for a user ID and password, SSO is not working because either the Domino or the WebSphere Application Server security server cannot authenticate the user with the Lightweight Directory Access Protocol (LDAP) server. Check the following possibilities:

- Verify that the LDAP server is accessible from the Domino server machine. Use the TCP/IP ping utility to check TCP/IP connectivity and to verify that the host machine is running.
- Verify that the LDAP user is defined in the LDAP directory. Use the **idsldapsearch** utility to confirm that the user ID exists and that the password is correct. For example, you can run the following command, entered as a single line:

You can use the OS/400® Qshell, a UNIX shell, or a Windows DOS prompt

```
% ldapsearch -D "cn=John Doe, ou=Rochester, o=IBM, c=US" -w mypassword  
-h myhost.mycompany.com -p 389 -b "ou=Rochester, o=IBM, c=US" (objectclass=*)
```

The percent character (%) indicates the prompt and is not part of the command. A list of directory entries is expected. Possible error conditions and causes are contained in the following list:

- No such object: This error indicates that the directory entry referenced by either the user's distinguished name (DN) value, which is specified after the **-D** option, or the base DN value, which is specified after the **-b** option, does not exist.
- Credentials that are not valid: This error indicates that the password is not valid.
- Cannot contact the LDAP server: This error indicates that the host name or the port specified for the server is not valid or that the LDAP server is not running.
- An empty list means that the base directory that is specified by the **-b** option does not contain any directory entries.
- If you are using the user's short name or user ID instead of the distinguished name, verify that the directory entry is configured with the short name. For a Domino directory, verify the **Short name/UserID** field of the Person document. For other LDAP directories, verify the **userid** property of the directory entry.

- If Domino authentication fails when using an LDAP directory other than a Domino directory, verify the configuration settings of the LDAP server in the Directory assistance document in the Directory assistance database. Also verify that the Server document refers to the correct Directory assistance document. The following LDAP values that are specified in the Directory Assistance document must match the values specified for the user registry in the WebSphere Application Server administrative domain:
 - Domain name
 - LDAP host name
 - LDAP port
 - Base DN

Additionally, the rules that are defined in the Directory assistance document must refer to the base distinguished name (DN) of the directory that contains the directory entries of the users.

You can trace Domino server requests to the LDAP server by adding the following line to the server `notes.ini` file:

```
webauth_verbose_trace=1
```

After restarting the Domino server, trace messages are displayed in the Domino server console as Web users attempt to authenticate to the Domino server.

- Authorization failure when accessing a protected resource.

After authenticating successfully, if an authorization error message is displayed, security is not configured correctly. Check the following possibilities:

- For Domino databases, verify that the user is defined in the access-control settings for the database. Refer to the Domino administrative documentation for the correct way to specify the user's DN. For example, for the DN `cn=John Doe, ou=Rochester, o=IBM, c=US`, the value on the access-control list must be set as `John Doe/Rochester/IBM/US`.
- For resources that are protected by WebSphere Application Server, verify that the security permissions are set correctly.
 - If granting permissions to selected groups, make sure that the user attempting to access the resource is a member of the group. For example, you can verify the members of the groups by using the following Web site to display the directory contents: `Ldap://myhost.mycompany.com:389/ou=Rochester, o=IBM, c=US??sub`
 - If you changed the LDAP configuration information (host, port, and base DN) in a WebSphere Application Server administrative domain since the permissions were set, the existing permissions are probably not valid and need to be recreated.

- SSO failure when accessing protected resources.

If a Web user is prompted to authenticate with each resource, SSO is not configured correctly. Check the following possibilities:

1. Configure both WebSphere Application Server and the Domino server to use the same LDAP directory. The HTTP cookie that is used for SSO stores the full DN of the user, for example, `cn=John Doe, ou=Rochester, o=IBM, c=US`, and the domain name service (DNS) domain.
2. Define Web users by hierarchical names if the Domino directory is used. For example, update the **User name** field in the Person document to include names of this format as the first value: `John Doe/Rochester/IBM/US`.
3. Specify the full DNS server name, not just the host name or TCP/IP address for Web sites issued to Domino servers and WebSphere Application Servers that are configured for SSO. For browsers to send cookies to a group of servers, the DNS domain must be included in the cookie, and the DNS domain in the cookie must match the Web address. This requirement is why you cannot use cookies across TCP/IP domains.
4. Configure both Domino and the WebSphere Application Server to use the same DNS domain. Verify that the DNS domain value is exactly the same, including capitalization. You need the name of the DNS domain in which WebSphere Application Server is configured. For additional information about configuring DNS domains for SSO, refer to the Single sign-on topic in the *Securing applications and their environment* PDF book.

5. Verify that the clustered Domino servers have the host name populated with the full DNS server name in the server document. By using the full DNS server name, Domino Internet Cluster Manager (ICM) can redirect to cluster members using SSO. If this field is not populated, by default, ICM redirects Web addresses to clustered Web servers by using the host name of the server only. ICM cannot send the SSO cookie because the DNS domain is not included in the Web address. To correct the problem:
 - a. Edit the Server document.
 - b. Click **Internet Protocols > HTTP** tab.
 - c. Enter the full DNS name of the server in the **Host names** field.
 6. If a port value for an LDAP server is specified for a WebSphere Application Server administrative domain, edit the Domino Web SSO configuration document and insert a backslash character (\) into the value of the **LDAP Realm** field before the colon character (:). For example, replace `myhost.mycompany.com:389` with `myhost.mycompany.com\:389`.
- Users are not logged out after the HTTP session timer expires.
 If users of WebSphere Application Server log onto an application and sit idle longer than the specified HTTP session timeout value, the user information is not invalidated and user credentials stay active until LTPA token timeout occurs.
 After you apply PK25740, complete the following steps to log out users from the application after the HTTP session has expired.
 1. In the administrative console, click **Security > Global security**.
 2. Under Custom properties, click **New**.
 3. In the Name field, enter `com.ibm.ws.security.web.logoutOnHTTPSessionExpire`.
 4. In the Values field, enter `true`.
 5. Click Apply and Save to save the changes to your configuration.
 6. Resynchronize and restart the server.

Security authorization provider troubleshooting tips

This article describes the issues you might encounter using a Java Authorization Contract for Containers (JACC) authorization provider. Tivoli Access Manager is bundled with WebSphere Application Server as an authorization provider. However, you also can plug in your own authorization provider.

Tivoli Access Manager as a Java Authorization Contract for Containers authorization provider

You might encounter the following issues when using Tivoli Access Manager as a JACC authorization provider:

- The configuration of JACC might fail.
- The server might fail to start after configuring JACC.
- The application might not deploy properly.
- The `startServer` command might fail after you have configured Tivoli Access Manager or a clean uninstall did not take place after unconfiguring JACC.
- An "HPDIA0202w An unknown user name was presented to Access Manager" error might occur.
- An "HPDAC0778E The specified user's account is set to invalid" error might occur.
- An WASX7017E: Exception received while running file "InsuranceServicesSingle.jacl" error might occur.
- "Access denied exceptions accessing applications when using JACC" on page 1326
- "An "HPDBA0219E: An error occurred reading data from an SSL connection" might occur" on page 1327

External providers for Java Authorization Contract for Containers authorization provider

You might encounter the following issues when you use an external provider for JACC authorization:

- An "HPDJA0506E Invalid argument: Null or zero-length user name field for the ACL entry" error might occur.

The configuration of JACC might fail

If you have problems configuring JACC, check the following items:

- Ensure that the parameters are correct. For example, you do not want a number after TAM_Policy_server_hostname:7135, but you do want be a number after TAM_Authorization_server_hostname:7136 (for example, TAM_Authorization_server_hostname:7136:1).
- If a message such as "server can't be contacted" is displayed, it is possible that the host names or port numbers of the Tivoli Access Manager servers are incorrect, or that the Tivoli Access Manager servers have not started.
- Ensure that the password for the sec_master user is correct.
- Check the SystemOut.log file and search for the AMAS string to see if any error messages are present.

The server might fail to start after configuring JACC

If the server does not start after JACC is configured, check the following items:

- Ensure that WebSphere Application Server and Tivoli Access Manager use the same Lightweight Directory Access Protocol (LDAP) server.
- If the message "Policy Director Authentication failed" is displayed, ensure that the:
 - WebSphere Application Server LDAP server ID is the same as the "Administrator user" in the Tivoli Access Manager JACC configuration panel.
 - Verify that the Tivoli Access Manager Administrator distinguished name (DN) is correct.
 - Verify that the password of the Tivoli Access Manager administrator has not expired and is valid.
 - Ensure that the account is valid for the Tivoli Access Manager administrator.
- If a message such as socket can't be opened for xxxx (where xxxx is a number) is displayed, take the following actions:
 1. Go to the *profile_root/etc/tam* directory.
 2. Change xxxx to an available port number in the *amwas.commonconfig.properties* file, and the *amwas*cellName_dmgr.properties* file if the deployment manager failed to start. If the node failed to start, change xxx to an available port number in the *amwas*cellName_nodeName_.properties* file. If the Application Server failed to start, change xxxx in the *amwas*cellname_nodeName_serverName.properties* file.

The application might not deploy properly

When you click **Save**, the policy and role information is propagated to the Tivoli Access Manager policy. This process might take some time to finish. If the save fails, you must uninstall the application and then reinstall it.

To access an application after it is installed, you must wait 30 seconds, by default, to start the application after you save.

The startServer command might fail after you configure Tivoli Access Manager or a clean uninstall did not take place after unconfiguring JACC.

If the cleanup for JACC unconfiguration or start server fails after JACC is configured, take the following actions:

- Remove Tivoli Access Manager properties files from WebSphere Application Server. For each application server in a Network Deployment (ND) environment with N servers defined (for example, server1, server2).

The following files must be removed.

```
install_root/java/jre/PdPerm.properties
install_root/java/jre/PdPerm.ks
profile_root/etc/tam/*
```

- Use a utility to clear the security configuration and return the system to the state it was in before you configure the JACC provider for Tivoli Access Manager. The utility removes all of the PDLoginModuleWrapper entries as well as the Tivoli Access Manager authorization table entry from the security.xml file, effectively removing the JACC provider for Tivoli Access Manager. Backup the security.xml file before running this utility.

Enter the following commands:

```
install_root/java/jre/bin/java -classpath
"install_root/lib/AMJACCProvider.jar:CLASSPATH"
com.tivoli.pdas.jacc.cfg.CleanSecXML fully_qualified_path/security.xml
```

An "HPDIA0202w An unknown user name was presented to Access Manager" error might occur

You might encounter the following error message if you try to use an existing user in a Local Directory Access Protocol (LDAP) user registry with Tivoli Access Manager:

```
AWXJR0008E Failed to create a PDPrincipal for principal mgr1.:
AWXJR0007E A Tivoli Access Manager exception was caught. Details are:
"HPDIA0202W An unknown user name was presented to Access Manager."
```

This problem might be caused by the host name exceeding predefined limits with Tivoli Access Manager when it is configured against MS Active Directory. In WebSphere Application Server, the maximum length of the host name can not exceed 46 characters.

Check that the host name is not fully qualified. Configure the machine so that the host name does not include the host domain.

To correct this error, complete the following steps:

1. On the command line, type the following information to get a Tivoli Access Manager command prompt:

```
pdadmin -a administrator_name -p administrator_password
```

The pdadmin *administrator_name* prompt is displayed. For example:

```
pdadmin -a administrator1 -p password
```

2. At the pdadmin command prompt, import the user from the LDAP user registry to Tivoli Access Manager by typing the following information:

```
user import user_name cn=user_name,o=organization_name,c=country
```

For example:

```
user import jstar cn=jstar,o=ibm,c=us
```

After importing the user to Tivoli Access Manager, you must use the **user modify** command to set the user account to valid. The following syntax shows how to use this command:

```
user modify user_name account-valid yes
```

For example:

```
user modify jstar account-valid yes
```

For information on how to import a group from LDAP to Tivoli Access Manager, see the Tivoli Access Manager documentation.

An "HPDAC0778E The specified user's account is set to invalid" error might occur

You might encounter the following error message after you import a user to Tivoli Access Manager and restart the client:

```
AWXJR0008E Failed to create a PDPrincipal for principal mgr1.:
AWXJR0007E A Tivoli Access Manager exception was caught.
Details are: "HPDAC0778E The specified user's account is set to invalid."
```

To correct this error, use the `user modify` command to set the user account to valid. The following syntax shows how to use this command:

```
user modify user_name account-valid yes
```

For example:

```
user modify jstar account-valid yes
```

An "HPDJA0506E Invalid argument: Null or zero-length user name field for the ACL entry" error might occur

You might encounter an error similar to the following message when you propagate the security policy information from the application to the provider using the `wsadmin propagatePolicyToJACCProvider` command:

```
AWXJR0035E An error occurred while attempting to add member,
           cn=agent3,o=ibm,c=us, to role AgentRole
HPDJA0506E Invalid argument: Null or zero-length user name field for
           the ACL entry
```

To correct this error, create or import the user, that is mapped to the security role to the Tivoli Access Manager. For more information on propagating the security policy information, see the documentation for your authorization provider.

An WASX7017E: Exception received while running file "InsuranceServicesSingle.jacl" error might occur

After the JACC provider and Tivoli Access Manager are enabled, when attempting to install the application, which is configured with security roles using the `wsadmin` command, the following error might occur:

```
WASX7017E: Exception received while running file "InsuranceServicesSingle.jacl";
exception information: com.ibm.ws.scripting.ScriptingException: WASX7111E:
Cannot find a match for supplied option:
"[RuleManager, , , cn=mgr3,o=ibm,c=us|cn=agent3,o=ibm,c=us, cn=ManagerGro
up,o=ibm,c=us|cn=AgentGroup,o=ibm,c=us]" for task "MapRolesToUsers"
```

The `$AdminApp MapRolesToUsers` task option is no longer valid when Tivoli Access Manager is used as the authorization server. To correct the error, change `MapRolesToUsers` to `TAMMapRolesToUsers`.

Access denied exceptions accessing applications when using JACC

In the case of Tivoli Access Manager, you might see the following error message.

```
AWXJR0044E: The access decision for Permission, {0}, was denied because either the
PolicyConfiguration or RoleConfiguration objects did not get created successfully at
application installation time. RoleConfiguration exists = {false}, PolicyConfiguration
exists = {false}."
```

If the access denied exceptions are not expected for the application, check the `SystemOut.log` files to see if the security policy information was correctly propagated to the provider.

If the security policy information for the application is successfully propagated to the provider, the audit statements with the message key SECJ0415I appear. However, if there was a problem propagating the security policy information to the provider (for example: network problems, JACC provider is not available), the SystemOut.log files contain the error message with the message keys SECJ0396E (during install) or SECJ0398E (during modification). The installation of the application is not stopped due to a failure to propagate the security policy to the JACC provider. Also, in the case of failure, no exception or error messages appear during the save operation. When the problem causing this failure is fixed, run the propagatePolicyToJaccProvider tool to propagate the security policy information to the provider without reinstalling the application. For more information about this task, see the Propagating security policy of installed applications to a JACC provider using wsadmin scripting topic in the *Securing applications and their environment* PDF book.

An "HPDBA0219E: An error occurred reading data from an SSL connection" might occur

An error message (HPDBA0219E) might appear in dmgr SystemOut.log when you install an application on WebSphere Application Server for Network Deployment (ND) and a managed node with Tivoli Access Manager is enabled.

If the error occurs, then the security policy data of recently deployed applications might not be immediately available. The policy data is available based on the server replicate time of the Tivoli Access Manager. This is defaulted to 30 seconds after all updates have been completed. To ensure that the latest policy data is available, log on to the pdadmin console and type: server replicate.

SPNEGO trust association interceptor (TAI) troubleshooting tips (deprecated)

Presented here is a list of trouble shooting tips useful in diagnosing Simple and Protected GSS-API Negotiation (SPNEGO) TAI problems and exceptions.

Note:

In WebSphere Application Server Version 6.1, a trust association interceptor (TAI) that uses the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) to securely negotiate and authenticate HTTP requests for secured resources was introduced. In WebSphere Application Server 7.0, this function is now deprecated. SPNEGO Web authentication has taken its place to provide dynamic reload of the SPNEGO filters and to enable fallback to the application login method.

The IBM Java Generic Security Service (JGSS) and IBM Simple and Protected GSS-API Negotiation (SPNEGO) providers use a Java virtual machine (JVM) custom property to control trace information. The SPNEGO TAI uses the JRas facility to allow an administrator to trace only specific classes. The following important trace specifications or JVM custom properties should be used to debug the TAI using tracing.

Table 66. SPNEGO TAI trace specifications

Trace	Use
<code>com.ibm.security.jgss.debug</code>	Set this JVM Custom Property to all to trace through JGSS code. Messages appear in the trace.log file, and SystemOut.log .
<code>com.ibm.security.krb5.Krb5Debug</code>	Set this JVM Custom Property to all to trace through the Kerberos5-specific JGSS code. Messages appear in the trace.log file, and SystemOut.log .
<code>com.ibm.ws.security.spnego.*</code>	Set this trace on using the administrative console > troubleshooting > Logging and Tracing > server1 > Change Log Detail Levels > com.ibm.ws.security.spnego.* . Messages appear in the trace.log file.

Problem: WebSphere Application Server and the Active Directory (AD) Domain Controller's time are not synchronized within 5 minutes.

Symptom

```
[2/24/06 13:12:46:093 CST] 00000060 Context      2 com.ibm.ws.security.spnego.Context
begin GSSContext accepted
[2/24/06 13:12:46:093 CST] 00000060 Context      E com.ibm.ws.security.spnego.Context
begin
CWSPN0011E: An invalid SPNEGO token has been encountered while authenticating a
HttpServletRequest:
0000: 60820160 06062b06 01050502 a1820154  ~..` ..+. .... ..T
0010: 30820150 a0030a01 01a10b06 092a8648  0..P .... .... *.H
0020: 82f71201 0202a282 013a0482 01366082  .... .... ..6`.
0030: 01320609 2a864886 f7120102 0203007e  .2.. *.H. .... ..~
0040: 82012130 82011da0 03020105 a1030201  ..!0 .... .... ....
0050: 1ea41118 0f323030 36303232 34313931  .... .200 6022 4191
0060: 3234365a a5050203 016b48a6 03020125  246Z .... .kH. ...%
0070: a9161b14 57535345 432e4155 5354494e  .... WSSE C.AU STIN
0080: 2e49424d 2e434f4d aa2d302b a0030201  .IBM .COM .-+ ....
0090: 00a12430 221b0448 5454501b 1a773230  ..$0 "...H TTP. .w20
00a0: 30337365 63646576 2e617573 74696e2e  03se cdev .aus tin.
00b0: 69626d2e 636f6dab 81aa1b81 a76f7267  ibm. com. .... .org
00c0: 2e696574 662e6a67 73732e47 53534578  .iet f.jg ss.G SSEX
00d0: 63657074 696f6e2c 206d616a 6f722063  cept ion, maj or c
00e0: 6f64653a 2031302c 206d696e 6f722063  ode: 10, min or c
00f0: 6f64653a 2033370a 096d616a 6f722073  ode: 37. .maj or s
0100: 7472696e 673a2044 65666563 74697665  trin g: D efec tive
0110: 20746f6b 656e0a09 6d696e6f 72207374  tok en.. mino r st
0120: 72696e67 3a20436c 69656e74 2074696d  ring : Cl ient tim
0130: 65204672 69646179 2c204665 62727561  e Fr iday , Fe brua
0140: 72792032 342c2032 30303620 61742031  ry 2 4, 2 006 at 1
0150: 3a31323a 34352050 4d20746f 6f20736b  :12: 45 P M to o sk
0160: 65776564  ewed
```

User Action

The preferred way to resolve this issue is to synchronize the WebSphere Application Server system time to within 5 minutes of the AD server's time. A best practice is to use a time server to keep all systems synchronized. You can also add or adjust the clockskew parameter in the Kerberos configuration file.

Note: The default for the clockskew parameter is 300 seconds (or 5 minutes).

Problem: No factory available to create a name for mechanism 1.3.6.1.5.5.2.

Problem

Getting an exception: No factory available to create a name for mechanism 1.3.6.1.5.5.2. There is no factory available to process the creation of a name for the specific mechanism.

Symptom

```
[4/8/05 22:51:24:542 EDT] 5003e481 SystemOut    0 [JGSS_DBG_PROV] Provider
IBMJGSSProvider version 1.01 does not support mech 1.3.6.1.5.5.2
[4/8/05 22:51:24:582 EDT] 5003e481 ServerCredent >
com.ibm.ws.security.spnego.ServerCredential initialize ENTRY
SPNEG0014: Kerberos initialization Failure: org.ietf.jgss.GSSEException, major code: 2,
minor code: 0
major string: Unsupported mechanism
minor string: No factory available to create name for mechanism 1.3.6.1.5.5.2
at com.ibm.security.jgss.i18n.I18NException.throwGSSEException
(I18NException.java:30)
at com.ibm.security.jgss.GSSManagerImpl.a(GSSManagerImpl.java:36)
at com.ibm.security.jgss.GSSCredentialImpl.add(GSSCredentialImpl.java:217)
at com.ibm.security.jgss.GSSCredentialImpl.<init>(GSSCredentialImpl.java:264)
```

User Action

Check the java.security file to ensure it contains the IBMSPNEGO security provider and that the provider is defined correctly. The java.security file should contain a line similar to:

```
security.provider.6=com.ibm.security.jgss.mech.spnego.IBMSPNEGO
```

Problem: Getting an exception as the JGSS library is trying to process the SPNEGO token.

Symptom

The following error is displayed as the JGSS library is trying to process the SPNEGO token.

```
Error authenticating request. Reporting to client
Major code = 11, Minor code = 31
org.ietf.jgss.GSSEException, major code: 11, minor code: 31
major string: General failure, unspecified at GSSAPI level
minor string: Kerberos error while decoding and verifying token:
com.ibm.security.krb5.internal.KrbException, status code: 31
message: Integrity check on decrypted field failed
```

Description This exception is the result of encoding the ticket using one key and attempting to decode it using a different key. There are number of possible reasons for this condition:

1. The Kerberos keytab file has not been copied to the server machine after it has been regenerated.
2. The Kerberos configuration points to the wrong Kerberos keytab file.
3. The Kerberos service principal name (SPN) has been defined to the Active Directory more than once. You have another userid defined with the same SPN or defined with the same SPN with a port defined also. The following example demonstrates how this condition can occur:

SAME SPN but different user ids

```
setspn -a HTTP/myHost.austin.ibm.com user1
setspn -a HTTP/myHost.austin.ibm.com user2
```

SAME SPN and same user ids, one without a port number, one with a port number

```
setspn -a HTTP/myHost.austin.ibm.com user
setspn -a HTTP/myHost.austin.ibm.com:9080 user
```

User Action If the problem is with the Kerberos keytab file, then regenerate the keytab file. If the problem is with multiple SPN definitions, then remove the extra or conflicting SPN, confirm that the SPN is no longer registered with the Active Directory, and then add the SPN. The Active Directory may need to be searched for other entries with SPNs defined that clash with the SPN.

To confirm that the SPN is not registered, the command:

```
setspn -l userid
```

should return with the following response:

```
Cannot find account userid
```

Problem: Single sign-on is not occurring.

Symptom When tracing is enabled, the following message appears:

```
[2/27/06 14:28:04:191 CST] 00000059 SpnegoHandler <
com.ibm.ws.security.spnego.SpnegoHandler handleRequest: Received a
non-SPNEGO Authorization Header RETURN
```

Description The client is returning an NT LAN manager (NTLM) response to the authorize challenge, not a SPNEGO token. This condition can be occur due to any of the following reasons:

- The client has not been configured properly.
- The client is not using a supported browser. For example, when using Microsoft Internet Explorer 5.5, SP1 responds with a non-SPNEGO authentication header.
- The user has not logged into the Active Directory domain, or into a trusted domain, or the client used does not support integrated authentication with Windows – in this case, the SPNEGO TAI is working properly.
- The user is accessing a service defined on the same machine upon which the client is running (local host). Microsoft Internet Explorer resolves the host name of the URL to `http://localhostsomeURL` instead of a fully qualified name.
- The SPN is not found in the Active Directory. The SPN must be of the format `HTTP/server.realm.com`. The command to add the SPN is

```
setspn -a HTTP/server.realm.com userid
```

- The Kerberos service principal name (SPN) has been defined to the Active Directory more than once. You have either another user ID defined with the same SPN or another userid defined with the same SPN with a port number defined. The following categories describe these conditions:

Same SPN but with differing user IDs

- `setspn -a HTTP/myappserver.austin.ibm.com user1`
- `setspn -a HTTP/myappserver.austin.ibm.com user2`

Same SPN and same user IDs, one with a port number defined

- `setspn -a HTTP/myappserver.austin.ibm.com user3`
- `setspn -a HTTP/myappserver.austin.ibm.com:9080 user3`

User Action

If the SPN is defined incorrectly as HTTP/server.realm.com@REALM.COM with the addition of @REALM.COM, then delete the user, redefine the user, and redefine the SPN.

If the problem is with the Kerberos keytab file, then regenerate the keytab file.

If the problem is with either category of multiple SPN definitions, then remove the extra or conflicting SPN, confirm that the SPN is no longer registered with the Active Directory, and then add the SPN. You can search the Active Directory for other SPN entries that are causing multiple SPN definitions. The following commands are useful to determine multiple SPN definitions:

```
setspn ?L userid
    Returns the message, cannot find account userid, if the SPN is not registered.

setspn -L
    Displays the SPNs that exist.
```

Problem: Credential Delegation is not working.**Symptom**

An invalid option is detected. When tracing is enabled, the following message is displayed:
com.ibm.security.krb5.KrbException, status code: 101 message: Invalid option in ticket request

Description

The Kerberos configuration file is not properly configured.

User Action

Ensure that neither renewable, nor proxiable are set to true.

Problem: Unable to get SSO working using RC4-HMAC encryption.**Symptom**

Examine the following message in the trace that you receive when trace is turned on:

```
com.ibm.security.krb5.internal.crypto.KrbCryptoException, status code: 0
message: Checksum error; received checksum does not match computed checksum
```

Description

RC4-HMAC encryption is not supported with a Microsoft Windows 2000 Kerberos key distribution center (KDC). To confirm this condition, examine the trace and identify where the exception is thrown. The content of the incoming ticket should be visible in the trace. Although the incoming ticket is encrypted, the SPN for the service is readable. If a Microsoft Windows 2000 KDC is used and the system is configured to use RC4-HMAC, the string representing the ticket for userid@REALM (instead of the expected HTTP/hostname.realm@REALM) is displayed. For example, this is beginning of the ticket received from a Microsoft Windows 2000 KDC:

```
0000: 01 00 6e 82 04 7f 30 82 04 7b a0 03 02 01 05 a1 ..n..0.....
0010: 03 02 01 0e a2 07 03 05 00 20 00 00 00 a3 82 03 .....
0020: a5 61 82 03 a1 30 82 03 9d a0 03 02 01 05 a1 0a .a..0.....
0030: 1b 08 45 50 46 44 2e 4e 45 54 a2 18 30 16 a0 03 ..REALM.COM.0..
0040: 02 01 01 a1 0f 30 0d 1b 0b 65 70 66 64 77 61 73 ....0...userid
0050: 75 6e 69 74 a3 82 03 6e 30 82 03 6a a0 03 02 01 .a.f...n0..j....
```

The realm is REALM.COM. The service name is userid. A correctly formed ticket for the same SPN is:

```
0000: 01 00 6e 82 04 56 30 82 04 52 a0 03 02 01 05 a1 ..n.V0..R.....
0010: 03 02 01 0e a2 07 03 05 00 20 00 00 00 a3 82 03 .....
0020: 82 61 82 03 7e 30 82 03 7a a0 03 02 01 05 a1 0a .a..0..z.....
0030: 1b 08 45 50 46 44 2e 4e 45 54 a2 2a 30 28 a0 03 ..REALM.COM.0...
0040: 02 01 02 a1 21 30 1f 1b 04 48 54 54 50 1b 17 75 ....0...HTTP..u
0050: 73 31 30 6b 65 70 66 77 61 73 73 30 31 2e 65 70 serid.realm.com.
0060: 66 64 2e 6e 65 74 a3 82 03 39 30 82 03 35 a0 03 ...n.....90..5..
```

User Action

To correct the problem, either use the Single data encryption standard (DES) or use a Microsoft Windows 2003 Server for a KDC. Remember to regenerate the SPN, and the Kerberos keytab file.

Problem: User receives the following message when accessing a protected URL through the SPNEGO SSO.**Symptom**

Examine the following message:

```
Bad Request
```

Your browser sent a request that this server could not understand.
Size of request header field exceeds server limit.

```
Authorization: Negotiate YII.....
```

Description	This message is generated by the Apache/IBM HTTP Server. This server is indicating that the authorization header returned by the user's browser is too large. The long string that follows the word Negotiate (in the error message above) is the SPNEGO token. This SPNEGO token is a wrapper of the Microsoft Windows Kerberos token. Microsoft Windows includes the user's PAC information in the Kerberos token. The more security groups that the user belongs to, the more PAC information is inserted in the Kerberos token, and the larger the SPNEGO becomes. IBM HTTP Server 2.0 (also Apache 2.0 and IBM HTTP Server 6.0) limit the size of any acceptable HTTP header to be 8K. In Microsoft Windows domains having many groups, and with user membership in many groups, the size of the user's SPNEGO token may exceed the 8K limit.
User Action	If possible, reduce the number of security groups the user is a member of. IBM HTTP Server 2.0.47 cumulative fix PK01070 allows for HTTP header sizes up to and beyond the Microsoft limit of 12K. WebSphere Application Server Version 6.0 users can obtain this fix in fixpack 6.0.0.2. Note: Non-Apache based Web servers may require differing solutions.

Problem: Even with JGSS tracing disabled, some KRB_DBG_KDC messages appear in the SystemOut.log.

Symptom	Examine the SystemOut.log and note the some KRB_DBG_KDC messages appear there even with JGSS tracing disabled.
Description	While most of the JGSS tracing is controlled by the com.ibm.security.jgss.debug property, a small set of messages are controlled by the com.ibm.security.krb5.Krb5Debug property. The com.ibm.security.krb5.Krb5Debug property has a default value to put some messages to the SystemOut.log
User Action	.To remove all KRB_DBG_KDC messages from the SystemOut.log , set the JVM property as follows: -Dcom.ibm.security.krb5.Krb5Debug=none

Problem: When an application contains a custom HTTP 401 error page, the SPNEGO TAI-generated HTTP response page is not displayed in a browser.

Symptom	When an application contains a custom HTTP 401 error page, the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) trust association interceptor (TAI)-generated HTTP response page is not displayed in a browser.
Description	When an application contains a custom HTTP 401 error page, the SPNEGO TAI-generated HTTP response page is not displayed in a browser. The custom HTTP 401 error page is displayed instead.
User Action	You can customize your HTTP 401 page to include information concerning how to configure your browser to use SPNEGO. For more information, see "Configuring the client browser to use SPNEGO TAI (deprecated)" on page 341 and "SPNEGO TAI custom properties configuration (deprecated)" on page 336.

Problem: HTTP Post parameters are lost during interaction with the SPNEGO TAI, when stepping down to userid/password login.

Symptom	Note that HTTP Post parameters are lost during interaction with the SPNEGO TAI, when stepping down to userid/password login.
Description	"Stepping down to userid/password login" means that the Microsoft Internet Explorer tries to respond initially with a SPNEGO token. If this response is unsuccessful, then the Microsoft Internet Explorer tries to respond with a NTLM token that is obtained through a userid/password challenge. The Microsoft Internet Explorer maintains state during a user's request. If a request was given the response of an "HTTP 401 Authenticate Negotiate", and the browser responds with a NTLM token obtained through a userid/password challenge, the browser resubmits the request. If this second request is given a response of an HTML page containing a redirection to the same URL but with new arguments (via Javascript) then the browser does not resubmit the POST parameters. Note: To avoid this problem, it is critical to NOT perform the automatic redirection. If the user clicks on a link, the problem does not occur.

User Action

The browser responds to the Authenticate/Negotiate challenge with an NTLM token, not an SPNEGO token. The SPNEGO TAI sees the NTLM, and returns back a HTTP 403 response, along with the HTML page. When the browser runs the Javascript `redirTimer` function, any POST or GET parameters that were present on the original request are lost.

By leveraging the `SPN<id>.NTLMTokenReceivedPage` property, an appropriate message page can be returned to the user. The default message that is returned (in the absence of a user defined property) is:

```
"<html><head><title>An NTLM Token was Received.</title></head>"
+ "<body>Your browser configuration is correct, but you have not logged into
  a supported Windows Domain."
+ "<p>Please login to the application using the normal login page.</html>";
```

Using the `SPN<id>.NTLMTokenReceivedPage` property, you can customize the exact response. It is critical that the returned HTML not perform a redirection.

When the SPNEGO TAI has been configured to use the shipped default `HTTPHeaderFilter` class as the `SPN<id>.filterClass`, then the `SPN<id>.filter` can be used to allow the second request to flow directly to the normal WebSphere Application Server security mechanism. In this way, the user experiences the normal authentication mechanism.

An example of such a configuration follows showing the required SPNEGO TAI properties necessary and the HTML file content.

***** SPNEGO TAI Property Name *****	***** HTML File Content *****
<code>com.ibm.ws.security.spnego.SPN1.hostName</code>	<code>server.wastedhed30.torolab.ibm.com</code>
<code>com.ibm.ws.security.spnego.SPN1.filterClass</code>	<code>com.ibm.ws.security.spnego.HTTPHeaderFilter</code>
<code>com.ibm.ws.security.spnego.SPN1.filter</code>	<code>request-ur!=noSPNEGO</code>
<code>com.ibm.ws.security.spnego.SPN1.NTLMTokenReceivedPage</code>	<code>File:///C:/temp/NTLM.html</code>

Note: Observe that the filter property instructs the SPNEGO TAI to NOT intercept any HTTP request that contains the string "noSPNEGO".

Here is an example of a generating a helpful response.

```
<html>
<head>
<title>NTLM Authentication Received </title>
<script language="javascript">
  var purl="" + document.location;
  if (purl.indexOf("noSPNEGO") < 0) {
    if (purl.indexOf('?') >= 0) purl += "&noSPNEGO";
    else purl += "?noSPNEGO";
  }
</script>
</head>
<body>
<p>An NTLM token was retrieved in response to the SPNEGO challenge. It is likely that
you are not logged into a Windows domain.<br>
Click on the following link to get the requested website.
<script language="javascript">
  document.write("<a href='"+purl+"'>");
  document.write("Open the same page using the normal authentication
  mechanism.");
  document.write("</a><br>");
</script>
You will not automatically be redirected.
</body>
</html>
```

Appendix. Directory conventions

References in product information to *app_server_root*, *profile_root*, and other directories infer specific default directory locations. This topic describes the conventions in use for WebSphere Application Server.

Default product locations (distributed)

The following file paths are default locations. You can install the product and other components or create profiles in any directory where you have write access. Multiple installations of WebSphere Application Server Network Deployment products or components require multiple locations. Default values for installation actions by root and non-root users are given. If no non-root values are specified, then the default directory values are applicable to both root and non-root users.

app_client_root

The following list shows default installation root directories for the WebSphere Application Client.

User	Directory
Root	<p>AIX /usr/IBM/WebSphere/AppClient (Java EE Application client only)</p> <p>HP-UX Linux Solaris /opt/IBM/WebSphere/AppClient (Java EE Application client only)</p> <p>Windows C:\Program Files\IBM\WebSphere\AppClient</p>
Non-root	<p>AIX HP-UX Linux Solaris <i>user_home</i>/IBM/WebSphere/AppServer/AppClient (Java EE Application client only)</p> <p>Windows C:\IBM\WebSphere\AppClient</p>

app_server_root

The following list shows the default installation directories for WebSphere Application Server Network Deployment.

User	Directory
Root	<p>AIX /usr/IBM/WebSphere/AppServer</p> <p>HP-UX Linux Solaris /opt/IBM/WebSphere/AppServer</p> <p>Windows C:\Program Files\IBM\WebSphere\AppServer</p>
Non-root	<p>AIX HP-UX Linux Solaris <i>user_home</i>/IBM/WebSphere/AppServer</p> <p>Windows C:\IBM\WebSphere\AppServer</p>

cip_app_server_root

A *customized installation package* (CIP) is an installation package created with IBM WebSphere Installation Factory that contains a WebSphere Application Server Network Deployment product bundled with one or more maintenance packages, an optional configuration archive, one or more optional enterprise archive files, and other optional files and scripts.

The following list shows the default installation root directories for a CIP where *cip_uid* is the CIP unique ID generated during creation of the build definition file.

User	Directory
Root	<p>AIX /usr/IBM/WebSphere/AppServer/cip/cip_uid</p> <p>HP-UX Linux Solaris /opt/IBM/WebSphere/AppServer/cip/cip_uid</p> <p>Windows C:\Program Files\IBM\WebSphere\AppServer\cip\cip_uid</p>
Non-root	<p>AIX HP-UX Linux Solaris user_home/IBM/WebSphere/AppServer/cip/cip_uid</p> <p>Windows C:\IBM\WebSphere\AppServer\cip\cip_uid</p>

component_root

The component installation root directory is any installation root directory described in this topic. Some programs are for use across multiple components. In particular, the Update Installer for WebSphere Software is for use with WebSphere Application Server Network Deployment, Web server plug-ins, the Application Client, and the IBM HTTP Server. All of these components are part of the product package.

gskit_root

IBM Global Security Kit (GSKit) can now be installed by any user. GSKit is installed locally inside the installing product's directory structure and is no longer installed in a global location on the target system. The following list shows the default installation root directory for Version 7 of the GSKit, where *product_root* is the root directory of the product that is installing GSKit, for example IBM HTTP Server or the Web server plug-in.

Directory
<p>AIX HP-UX Linux Solaris product_root/gsk7</p> <p>Windows product_root\gsk7</p>

if_root This directory represents the root directory of the IBM WebSphere Installation Factory. Because you can download and unpack the Installation Factory to any directory on the file system to which you have write access, this directory's location varies by user. IBM WebSphere Installation Factory is an Eclipse-based tool which creates installation packages for installing WebSphere Application Server in a reliable and repeatable way, tailored to your specific needs.

iip_root

This directory represents the root directory of an *integrated installation package* (IIP) produced by the IBM WebSphere Installation Factory. Because you can create and save an IIP to any directory on the file system to which you have write access, this directory's location varies by user. An IIP is an aggregated installation package that can include one or more generally available installation packages, one or more customized installation packages (CIPs), and other user-specified files and directories.

profile_root

The following list shows the default directory for a profile named *profile_name* on each distributed operating system.

User	Directory
Root	<p>AIX /usr/IBM/WebSphere/AppServer/profiles/profile_name</p> <p>HP-UX Linux Solaris /opt/IBM/WebSphere/AppServer/profiles/profile_name</p> <p>Windows C:\Program Files\IBM\WebSphere\AppServer\profiles\profile_name</p>

User	Directory
Non-root	<p>AIX HP-UX Linux Solaris</p> <p><code>user_home/IBM/WebSphere/AppServer/profiles/</code></p> <p>Windows <code>C:\IBM\WebSphere\AppServer\profiles\</code></p>

plugins_root

The following default installation root is for the Web server plug-ins for WebSphere Application Server.

User	Directory
Root	<p>AIX <code>/usr/IBM/WebSphere/Plugins</code></p> <p>HP-UX Linux Solaris <code>/opt/IBM/WebSphere/Plugins</code></p> <p>Windows <code>C:\Program Files\IBM\WebSphere\Plugins</code></p>
Non-root	<p>AIX HP-UX Linux Solaris</p> <p><code>user_home/IBM/WebSphere/Plugins</code></p> <p>Windows <code>C:\IBM\WebSphere\Plugins</code></p>

updi_root

The following list shows the default installation root directories for the Update Installer for WebSphere Software.

User	Directory
Root	<p>AIX <code>/usr/IBM/WebSphere/UpdateInstaller</code></p> <p>HP-UX Linux Solaris <code>/opt/IBM/WebSphere/UpdateInstaller</code></p> <p>Windows <code>C:\Program Files\IBM\WebSphere\UpdateInstaller</code></p>
Non-root	<p>AIX HP-UX Linux Solaris</p> <p><code>user_home/IBM/WebSphere/UpdateInstaller</code></p> <p>Windows <code>C:\IBM\WebSphere\UpdateInstaller</code></p>

web_server_root

The following default installation root directories are for the IBM HTTP Server.

User	Directory
Root	<p>AIX <code>/usr/IBM/HTTPServer</code></p> <p>HP-UX Linux Solaris <code>/opt/IBM/HTTPServer</code></p> <p>Windows <code>C:\Program Files\IBM\HTTPServer</code></p>
Non-root	<p>AIX HP-UX Linux Solaris</p> <p><code>user_home/IBM/HTTPServer</code></p> <p>Windows <code>C:\IBM\HTTPServer</code></p>

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Intellectual Property & Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

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 is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.