



Security Guide

/ ForgeRock Access Management 7.0.2

Latest update: 7.0.2

ForgeRock AS.
201 Mission St., Suite 2900
San Francisco, CA 94105, USA
+1 415-599-1100 (US)
www.forgerock.com

Copyright © 2019-2021 ForgeRock AS.

Abstract

Guide to securing the core of the ForgeRock Access Management deployment.



This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/3.0/> or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.

ForgeRock® and ForgeRock Identity Platform™ are trademarks of ForgeRock Inc. or its subsidiaries in the U.S. and in other countries. Trademarks are the property of their respective owners.

UNLESS OTHERWISE MUTUALLY AGREED BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

DejaVu Fonts

Bitstream Vera Fonts Copyright

Copyright (c) 2003 by Bitstream, Inc. All Rights Reserved. Bitstream Vera is a trademark of Bitstream, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Bitstream" or the word "Vera".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Bitstream Vera" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the names of Gnome, the Gnome Foundation, and Bitstream Inc., shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from the Gnome Foundation or Bitstream Inc., respectively. For further information, contact: fonts at gnome dot org.

Arev Fonts Copyright

Copyright (c) 2006 by Tavmjong Bah. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the modifications to the Bitstream Vera Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Tavmjong Bah" or the word "Arev".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Tavmjong Bah Arev" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL TAVMJONG BAH BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the name of Tavmjong Bah shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from Tavmjong Bah. For further information, contact: tavmjong @ free . fr.

FontAwesome Copyright

Copyright (c) 2017 by Dave Gandy, <https://fontawesome.com/>.

This Font Software is licensed under the SIL Open Font License, Version 1.1. See <https://opensource.org/licenses/OFL-1.1>.

Table of Contents

Overview	v
1. General Security Considerations	1
2. Securing Network Communication	3
Preventing Insecure HTTP and LDAP Connections	4
Configuring AM Behind a Reverse Proxy	8
Configuring CORS Support	15
Configuring the Cookie Domain	26
Cross-Site Request Forgery (CSRF) Protection	27
3. Securing Administrative Access	28
About the amAdmin User	28
Delegating Privileges	34
Securing the Administration Console	39
Securing Administrative Tools	39
4. Securing Realms	41
About the Demonstration User	42
5. Configuring Secrets, Certificates, and Keys	43
Features in AM That Use Keys	48
Managing the AM Keystore	50
Managing Key Aliases and Passwords	54
Configuring Secret Stores	62
Mapping and Rotating Secrets	73
Changing Default Key Aliases	82
6. Securing the Session Cookie	84
Configuring HttpOnly Session Cookies	85
Configuring Secure Session Cookies	85
Changing the name of the Session Cookie	86
Enabling Restricted Tokens for CDSSO Session Cookies	86
7. Additional Cookie Security Considerations	88
Enabling SameSite Cookie Rules	88
Managing the Secure Cookie Filter	89
Changing the Name of the Sticky Load Balancing Cookie	90
8. Securing Sessions	91
Understanding Session Termination	92
Configuring Account Lockout	95
Configuring Session Quotas	97
Configuring Client-Based Session Security	98
Protect Sensitive Attributes	102
Configuring Authentication Session Whitelisting	103
9. Request Security Considerations	104
Controlling the Maximum Size of Compressed JWTs	104
Limiting the Size of the Request Body	105
10. Protecting Applications	107
Identity Gateway or AM Web and Java Agents?	107
11. Setting Up Audit Logging	108

About the Audit Logging Service	109
Implementing the Audit Logging Service	110
Implementing the Classic Logging Service	132
12. Reference	134
Audit Logging Reference	134
Customizing CTS-Based Session Quota Exhaustion Actions	158
Glossary	163

Overview





This guide is written for administrators that are comfortable securing web applications. Although the authors will attempt to lay out a comprehensive list of actions to take, security is a too-broad subject to tackle and every environment is different; readers are expected to do their own research and complement the information found in this guide.

This guide will not provide guidance to secure advanced AM features, such as OAuth 2.0 or SAML v2.0. You will find the relevant information in their respective guides.

When you deploy AM, you must ensure that your environment is built and configured with security in mind. This includes:

- The network infrastructure.
- The operating system.
- The container where AM runs.
- The Java installation and the cryptography settings.
- The clients and applications that will connect to AM.
- The CTS store, identity stores, and any other application stores.
- AM's own configuration.

Quick Start

 Network Learn tips and best practices about securing your network infrastructure.	 Audit Logging Learn how AM logs auditing events and configure the Audit Logging Service to suit your needs.
 Certificates and keys Learn about cryptographic keys, keystores, and secret stores.	 Protect the Session Cookie from Hijacking Discover how to protect the session cookie from malicious users.

About ForgeRock Identity Platform™ Software

ForgeRock Identity Platform™ serves as the basis for our simple and comprehensive Identity and Access Management solution. We help our customers deepen their relationships with their customers, and improve the productivity and connectivity of their employees and partners. For more information about ForgeRock and about the platform, see <https://www.forgerock.com>.

Chapter 1

General Security Considerations

This list does not intend to show you best practices about network and system administration, but intends to make a number of points related to security that you can expand upon.

Keep up to date on patches

Security vulnerabilities are the reason why you should keep your operating systems, web and application servers, and any other application in your environment up to date. Knowledge of vulnerabilities spread fast across malicious users, who would not hesitate in trying to exploit them.

ForgeRock maintains a list of [security advisories](#) that you should follow. You should also follow similar lists from all your vendors.

Keep up to date on cryptographic methods and algorithms

Different algorithms and methods are discovered and tested over time, and communities of experts decide which are the most secure for different uses. Do not use outdated algorithms such as RSA for generating your keys.

Turn off unnecessary features

The more features you have turned on, the more features you need to secure, patch, and audit. If something is not being used, disable it or uninstall it.

Limit access to the servers hosting AM

A large part of protecting your environment is making sure only authorized people can access your servers and applications through the appropriate network, using the appropriate ports, and presenting strong-enough credentials.

Ensure users connect through SSL / TLS to the systems and audit system access periodically.

For a list of ports used in AM by default, see "[Ports Used](#)" in the *Reference*.

Enforce security

Do not expect your users to follow security practices on their own; enforce security when possible by requiring secure connections, password resets, and strong authentication methods.

Audit Access and Changes

Audit logs record all events that have happened. Some applications store them with their engine logs, some others use specific files or send the information to a different server for archiving.

Operating systems have audit logs as well, to detect unauthorized login attempts and changes to the software.

AM has its own audit logging service that adheres to the log structure common across the ForgeRock Identity Platform.

Chapter 2

Securing Network Communication

Keeping your AM instances safe from both internal and external attacks is paramount, but it is also a challenge when you cannot control who connects to your instances.

For example, a client could send unprotected credentials in an HTTP Authorization header. Even if AM were to reject the request, the credentials would already be leaked to any eavesdroppers.

The best way to protect your environment is to enforce the use of secure HTTPS communication.

The following table summarizes best practices about network security in AM environments:

Task	Resources
Enforce Secure Connections Secure connections between AM and the rest of your platform, whether it is DS servers or your applications.	"Preventing Insecure HTTP and LDAP Connections"
Use a Reverse Proxy Configure AM behind a reverse proxy. This will protect AM against DoS attacks and restrict access to AM and its endpoints to networks you trust.	"Configuring AM Behind a Reverse Proxy"
Configure CORS filters Configure a CORS filter such that only your trusted clients and applications can make cross-domain calls to your AM instances.	"Configuring CORS Support"
Adjust AM's Cookie Domain Configure AM cookie domain so that AM communicates with the hosts in the required domains and sub-domains.	"Configuring the Cookie Domain"
Learn about the CSRF Protection Filter for REST endpoints By default, AM protects its <code>/json</code> endpoints using a header filter.	"Cross-Site Request Forgery (CSRF) Protection"

Preventing Insecure HTTP and LDAP Connections

Both HTTPS and LDAPS secure connections are based on the transport layer security protocol (TLS), which depends on digital certificates (also called public key certificates).

Digital certificates are for sharing public keys used for signing and encryption, and they include information such as the public key, the owner of such key, and a digital signature created by the issuer of the certificate.

In client-server environments, the server provides a certificate that proves that the content it serves is as intended and has not been modified by malicious users. In some environments, however, the client is also required to present its own certificate; this is what is called mutual TLS (or mTLS).

In order to begin the TLS handshake, the actor receiving the certificate must know and trust the issuer of the certificate. This happens by default for certificates issued by a certificate authority (CA), but never for self-signed certificates. This means that, if you decide to have self-signed certificates, you must share them across the servers and applications that need to communicate in your environment.

Tip

Be mindful of security breaches and vulnerabilities that happen across the world, and ensure your environment is not using outdated insecure protocols, such as SSL 3.0, TLS 1.0, and others.

Configuring the AM Container for HTTPS Connections

Configure the container where AM runs for HTTPS to prevent communication over insecure HTTP. This includes HTTPS communication between AM and Web/Java Agents, and AM and your applications, or AM and any other member of the ForgeRock Identity Platform.

Note that configuring AM for HTTPS is the first step; you need to also configure the Web/Java agent, your applications, and any other member of the ForgeRock Identity Platform for HTTPS, too.

HTTPS connections happen at container level, encapsulated in the TLS protocol. This means AM itself is not involved in checking or sending certificates. The same is true for Web and Java agents.

Some advanced AM features, however, require AM to be able to validate certificates without the mediation of the container. For more information about those features, see *"Configuring Secrets, Certificates, and Keys"*.

To secure communications to AM, configure the container for HTTPS connections and install AM using the **https** protocol and the appropriate secure port. Follow the steps in [Installation Guide](#) to prepare your environment and install AM.

You can also reconfigure your instances to use HTTPS. For more information, see *How do I enable SSL in AM/OpenAM (All versions) for an existing installation?* in the Knowledge Base.

To control the protocols used for outbound HTTPS connections, configure the `-Dhttps.protocols` JVM setting in the container where AM runs. For details, see "Security Settings" in the *Installation Guide*.

Securing Directory Server Communication

Configure AM and the identity and data stores that connect to it to enforce secure communication, either using LDAPS or StartTLS. This includes communication between AM and the CTS store, between AM and the application stores, and between AM and the identity stores.

To Configure AM to Trust Directory Server Certificates

Secure directory server connections check certificates stored in the truststore of the container where AM runs. This procedure assumes you are using Apache Tomcat and a DS instance. Refer to your container and directory server documentation for more information.

1. Configure your stores to enforce secure communication, if they do not already.

For DS instances, see *Require LDAPS* in the *ForgeRock Directory Services Security Guide*.

Note

DS 7 or later is configured to require secure connections by default; therefore, you might have already configured some of your stores to use secure connections during the AM installation process.

2. Export the DS server certificate:

```
$ /path/to/openssl/bin/dskeymgr export-ca-cert \  
--deploymentKey $DEPLOYMENT_KEY \  
--deploymentKeyPassword password \  
--alias ds-ca-cert \  
--outputFile ds-ca-cert.pem
```

Note that `$DEPLOYMENT_KEY` is a Unix variable that contains the DS deployment key, so that it is not logged in the user's command history.

The default DS server certificate only has the hostname you supplied at setup time, and `localhost`, as the value of the `SubjectAlternativeName` attribute; however, certificate hostname validation is strict.

Ensure that the certificate matches the hostname (or the FQDN) of the DS server before continuing.

3. Import the DS certificate into the AM truststore:

```
$ keytool \  
-importcert \  
-alias ds-ca-cert \  
-file ds-ca-cert.pem \  
-keystore /path/to/openam/security/keystores/truststore
```

You are now ready to configure AM to use secure connections to the directory server.

To Secure Directory Server Communication

1. Make a backup of your environment, as explained in "[Backing Up Configurations](#)" in the *Maintenance Guide*.
2. Ensure your stores are ready for secure connections, and that AM can trust the certificates of the directory servers. Failure to do so may cause several problems, such as the `amAdmin` user being unable to log in, or AM being unable to start up.

Try the change first in test or development environments.

Tip

Certificate hostname validation is strict. AM checks that the hostname in the LDAP server certificate matches the hostname of the directory server, and DS checks that the server it is trying to connect to has a certificate that matches its hostname.

3. Specify the TLS protocol(s) AM will use for outbound LDAPS connections by configuring the `-Dorg.forgerock.openam.ldap.secure.protocol.version` JVM setting in the container where AM runs. For example:

```
-Dorg.forgerock.openam.ldap.secure.protocol.version=TLSv1.2,TLSv1.3
```

For details, see "[Security Settings](#)" in the *Installation Guide*

4. Log in to the AM console with an administrative user, for example, `amAdmin`.
5. To configure identity stores:
 - a. Navigate to Realms > *Realm Name* > Identity Stores > *Store Name* > Server Settings.
 - b. In the LDAP Connection Mode drop-down list, select LDAPS.
 - c. Save your changes.

Perform these steps on every realm as necessary.

6. To configure LDAPS for the external CTS store:
 - a. Navigate to Deployment > Servers > *Server Name* > CTS > External Store Configuration.
 - b. Enable the SSL/TLS Enabled switch.
 - c. Save your changes.
7. To configure the configuration store:

- a. Navigate to Deployment > Servers > *Server Name* > Directory Configuration > Server.
- b. On the Connection type dropdown list, select **SSL**.
- c. Save your changes.

Perform these steps on every server as necessary.

8. To configure external policy and application stores:

- a. Navigate to Configure > Global Service > External Data Stores > Secondary Configurations > *Store Name*.
- b. Enable the Use SSL switch.
- c. Save your changes.

Perform these steps for each store on every realm as necessary.

9. To configure external UMA stores:

- a. Navigate to Deployment > Servers > *Server Name* > UMA > *External UMA store*.
- b. Enable the SSL/TLS Enabled switch.
- c. Save your changes.

Perform these steps for each store as necessary.

10. When using clients, ensure you make LDAP calls through the LDAPS port and that the client has access to the store certificate. Otherwise, the LDAP server will not be able to validate the connection.

For DS stores, you should also specify the keystore file containing the store certificate, and its password. For example:

```
--port 1636 \  
--useSsl \  
--usePkcs12TrustStore /path/to/openssl/config/keystore \  
--trustStorePasswordFile /path/to/openssl/config/keystore.pin \  

```

Different commands may require different options. Different keystore types, too. For more information, see the *Directory Services Tools Reference* .

Configuring AM Behind a Reverse Proxy

Reverse proxies (such as ForgeRock Identity Gateway) are proxy servers that sit between clients and application servers. Their main function is to act on behalf of the application server, forwarding resources to the client as if they were the application server itself.

Modern reverse proxies also provide additional functionality such as load balancing, compression, SSL termination, web acceleration, and firewall capabilities.

Configuring a reverse proxy in front of your AM instances provides the following security benefits:

- Protecting AM servers from denial of service attacks.

A reverse proxy will terminate incoming connections and reopen them against the AM servers, effectively masking the {`$am.abbr`} IP addresses. This makes it more difficult for attackers to launch DoS attacks against them. A firewall can prevent direct access to the AM servers.

- SSL termination/SSL offloading.

Since reverse proxies terminate incoming connections to AM, they also decrypt the HTTPS requests and pass them unencrypted to the container where AM runs.

This has several benefits, such as removing the need to install certificates in the containers, which simplifies the management of SSL/TLS.

Depending on your environment, though, you may decide to configure SSL/TLS between AM and the reverse proxy, or configure the proxy to pass-through the SSL traffic to the container where AM runs.

This guide, and the examples in other AM guides default to AM being configured to use HTTPS communication.

- Provide a unique point of access to AM.

Configuring a reverse proxy in front of AM creates a channel between the public network and the internal network.

Since all communication to AM needs to come from the reverse proxy, you can, for example, restrict access to a set of trusted networks. You can fine-tune the access restrictions for each request and apply rate-limiting and load balancing such that a possible attack does not bring down your whole infrastructure.

- Protect Endpoints

In the same way that you can restrict access to trusted networks, you can also restrict access to any endpoint AM is exposing.

AM exposes a number of internal administration endpoints, such as the `/sessionservice` endpoint. You must ensure those are not reachable over the Internet.

For a list of internal endpoints that you should protect, see "*Service Endpoints*" in the *Reference*.

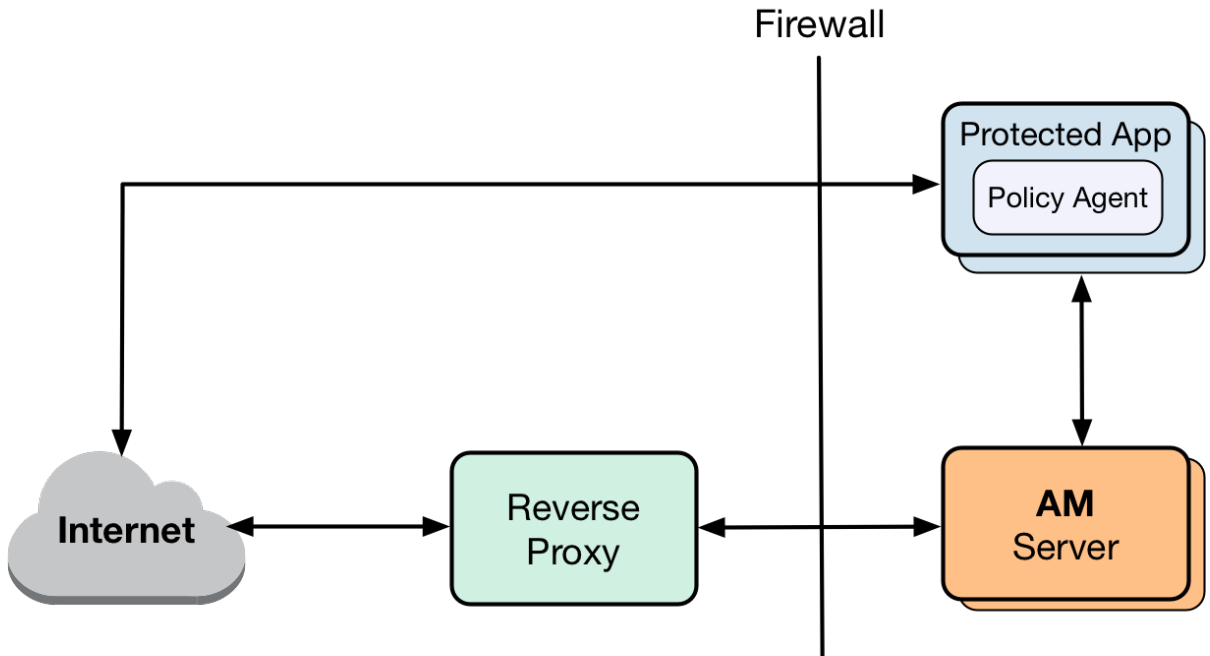
Regarding feature endpoints, AM makes endpoints accessible the moment an administrator creates a service. For example, the OAuth 2.0 endpoints are not available by default, but configuring an instance of the OAuth 2.0 provider service in a realm will make the endpoints available for that realm.

You must ensure you are exposing the correct endpoints to the Internet.

Recommending how to set up your network infrastructure is beyond the scope of this guide. There are too many permutations that are valid use cases; for example, some environments may deploy a reverse proxy for its load balancing capabilities instead of dedicated, hardware-based load balancers. More complex deployments may have multiple layers of firewalls, load balancers, and reverse proxies.

The following figure is an example of a possible configuration:

Exposing Only a Reverse Proxy to the Internet



Architecture protecting your AM services behind an Internet-facing reverse proxy

The following table summarizes the high-level tasks required to configure AM when it is behind a proxy:

Task	Resources
Configure the Proxy's Details Configure AM or the container where it runs to route outbound traffic through the proxy.	"Configuring AM for Outbound Communication"
Configure the Base URL Service	"Configuring the Base URL Source Service"

Task	Resources
Services configure their endpoints based on AM's URL. The Base URL Service remaps the endpoints of the services that require it to the proxy's URL.	

Configuring AM for Outbound Communication

Clients from different networks connect to AM to use its functionality. These clients initiate communication with AM and the container where it runs. However, when AM acts as a client to a third-party application, it makes outbound calls outside its container to retrieve information or services.

When AM is behind a proxy, you must route AM's client through the proxy. To do so, provide the proxy's details to AM and the container where it runs:

1. Set the relevant proxy JVM options in the container where AM runs.

+ HTTPS Options

-Dhttps.proxyHost

IP address or hostname of the proxy server. For example, `proxy.example.com`.

-Dhttps.proxyPort

Port number of the proxy server. For example, `8443`.

-Dhttp.nonProxyHosts

A pipe-separated (|) list of IP addresses or hostnames that should be reached directly, bypassing the proxy configuration. For example, `localhost|internal.example.com`.

Use wildcards (*) at the beginning or the end of the address or hostname. For example, `*.example.com` or `internal*`.

+ HTTP Options

-Dhttp.proxyHost

IP address or hostname of the proxy server. For example, `proxy.example.com`.

-Dhttp.proxyPort

Port number of the proxy server. For example, `8080`.

-Dhttp.nonProxyHosts

A pipe-separated (|) list of IP addresses or hostnames that should be reached directly, bypassing the proxy configuration. For example, `localhost|internal.example.com`.

Use wildcards (*) at the beginning or the end of the address or hostname. For example, `*.example.com` or `internal*`.

For example, set the properties in the `JAVA_OPTS` variable of the `$CATALINA_BASE/bin/setenv.sh` Apache Tomcat file.

2. Understand whether your proxy requires authentication:

- If the proxy requires authentication:
 - a. In the `org.forgerock.openam.httpclienthandler.system.proxy.uri` advanced server property, configure the URI of the proxy. The URI must be in the format `scheme://hostname:port`. For example, `https://myproxy.example.com:443`.

+ How Do I Configure Advanced Server Properties?

- To configure advanced server properties in the AM Admin UI for all AM instances, go to Configure > Server Defaults > Advanced.
- To configure advanced server properties for a particular instance, go to Deployment > Servers > *Server Name* > Advanced.

If the property you want to add or edit is not already configured, add it with its value, then click on the plus (+) button.

If the property you want to add or edit is already configured, click on the pencil (✎) button to edit it. When you are finished, click on the tick (✓) button.

Save your changes.

- b. In the `org.forgerock.openam.httpclienthandler.system.proxy.username` and the `org.forgerock.openam.httpclienthandler.system.proxy.password` advanced server properties, configure the proxy's credentials.
- c. In the `org.forgerock.openam.httpclienthandler.system.nonProxyHosts` advanced server property, provide one or more target hosts for which resulting HTTP client requests should *not* be proxied.

The list must comma-separated, for example `localhost,127.*,*.example.com`.

Configuring these properties lets features using ForgeRock's ClientHandler code use the proxy settings defined in the advanced server properties.

- If the proxy does not require authentication:
 - Set the `org.forgerock.openam.httpclienthandler.system.proxy.enabled` advanced server property to `true`.

+ *How Do I Configure Advanced Server Properties?*

- To configure advanced server properties in the AM Admin UI for all AM instances, go to Configure > Server Defaults > Advanced.
- To configure advanced server properties for a particular instance, go to Deployment > Servers > *Server Name* > Advanced.

If the property you want to add or edit is not already configured, add it with its value, then click on the plus (+) button.

If the property you want to add or edit is already configured, click on the pencil (✎) button to edit it. When you are finished, click on the tick (✓) button.

Save your changes.

Configuring this property lets features using ForgeRock's ClientHandler code use the JVM proxy settings.

For more information about the advanced server properties, see "Advanced Properties" in the *Reference*.

Tip

You can tune the connection factory behavior of the features that use ForgeRock's ClientHandler code. For example, the scripting engine, or the social provider authentication nodes.

+ *Client Connection Handler Properties*

The following advanced server properties control different aspects of the connection factory:

- `org.forgerock.openam.httpclienthandler.system.clients.connection.timeout`
- `org.forgerock.openam.httpclienthandler.system.clients.max.connections`
- `org.forgerock.openam.httpclienthandler.system.clients.pool.ttl`
- `org.forgerock.openam.httpclienthandler.system.clients.response.timeout`
- `org.forgerock.openam.httpclienthandler.system.clients.retry.failed.requests.enabled`

- `org.forgerock.openam.httpclienthandler.system.clients.reuse.connections.enabled`

They have sensible defaults configured, but if you need to change them, see Advanced Properties.

Configuring the Base URL Source Service

In many deployments, AM determines the base URL of a provider using the incoming HTTP request. However, there are often cases when the base URL of a provider cannot be determined from the incoming request alone, especially if the provider is behind some proxying application. For example, if an AM instance is part of a site where the external connection is over SSL but the request to the AM instance is over plain HTTP, then AM would have difficulty in reconstructing the base URL of the provider.

In these cases, AM supports a provider service that allows a realm to have a configured option for obtaining the base URL including protocol for components that need to return a URL to the client.

To Configure the Base URL Source Service

1. Log in to the AM console as an administrative user, such as `amAdmin`, and then navigate to Realms > *Realm Name* > Services.
2. Click Add a Service, select Base URL Source, and then click Create, leaving the fields empty.
3. For Base URL Source, select one of the following options:

Base URL Source Options

Option	Description
Extension class	Click the Extension class to return a base URL from a provided <code>HttpServletRequest</code> object. In the Extension class name field, enter <code>org.forgerock.openam.services.baseurl.BaseURLProvider</code> .
Fixed value	Click Fixed value to enter a specific base URL value. In the Fixed value base URL field, enter the base URL.
Forwarded header	Click Forwarded header to retrieve the base URL from the <code>Forwarded</code> header field in the HTTP request. The Forwarded HTTP header field is standardized and specified in <i>RFC 7239</i> .
Host/protocol from incoming request (default)	Click Host/protocol from incoming request to get the hostname, server name, and port from the HTTP request.
X-Forwarded-* headers	Click X-Forwarded-* headers to use non-standard header fields, such as <code>X-Forwarded-For</code> , <code>X-Forwarded-By</code> , and <code>X-Forwarded-Proto</code> .

4. In the Context path, enter the context path for the base URL. If provided, the base URL includes the deployment context path appended to the calculated URL. For example, `/openam`.
5. Click Finish to save your configuration.

Configuring CORS Support

Cross-origin resource sharing (CORS) allows requests to be made across domains from user agents.

To configure CORS support in AM, use the global CORS service UI, or use the `/global-config/services/CorsService` REST endpoint.

The configurations you create with either method are combined to form the entire set of rules for resource sharing. The CORS service also collects the values of the JavaScript Origins property in each OAuth 2.0 client configured, and adds them to the list of accepted origins.

Any changes you make to CORS configurations, using either the UI or REST, take effect immediately without requiring a restart.

Note

In previous AM releases, you configured CORS filters in the deployment descriptor file (`web.xml`). This method of configuring CORS is not supported, from AM version 7 onwards.

Configure CORS (UI)

You can use the UI to add multiple CORS configurations to AM, which are combined and used to ensure that only your trusted clients and applications can access your AM instance's resources.

For example, you could use the REST endpoint to add a base configuration, allowing a broad set of headers, and then add a stricter configuration; for example, for your OAuth 2.0 clients.

- "To Enable the CORS Filter"
- "To Add a CORS Configuration"
- "To Delete a CORS Configuration"

To Enable the CORS Filter

To enable CORS globally, go to Configure > Global Services > CORS Service > Configuration, and enable the Enable the CORS filter property.

If this property is not enabled, no CORS headers are added to any responses from AM, and CORS is disabled.

To Add a CORS Configuration

To add a CORS configuration, go to Configure > Global Services > CORS Service > Secondary Configurations, and then click Add a Secondary Configuration

The initial page that appears contains the following properties:

Name

Provide a descriptive name for the configuration to make management of multiple rules easier.

Accepted Origins

Add the *origins* allowed when making CORS requests to AM. Wildcards are not supported; each value should be an exact match for the origin of the CORS request.

The CORS service automatically collects the values of the JavaScript Origins property in each OAuth 2.0 client configured, and adds them to an internal list of accepted origins. You do not need to add them manually.

Tip

During development you may not be using FQDNs as the origin of a CORS request; for example, when you are using the `file://` protocol locally.

If so, you can add these non-FQDN origins to the list; for example, `file://` and `null`.

Accepted Methods

Add the HTTP methods allowed when making CORS requests to AM. The list is included in pre-flight responses, in the `Access-Control-Allow-Methods` header.

The method names are case-sensitive, ensure they are entered in all uppercase characters.

Accepted Headers

Add the request header names allowed when making CORS requests to AM. The list is included in pre-flight responses, in the `Access-Control-Allow-Headers` header.

The header names are case-insensitive.

By default, the following simple headers are explicitly accepted:

- `Cache-Control`
- `Content-Language`
- `Expires`
- `Last-Modified`
- `Pragma`

If you do not specify values for this element, the presence of any header in the CORS request, other than the simple headers listed above, will cause the request to be rejected.

+ *What are the commonly used headers?*

Headers commonly used when accessing an AM server include the following:

Commonly Used Headers

Header	Information
<code>iPlanetDirectoryPro</code>	Used for session information. See <i>"Introducing Sessions"</i> in the <i>Sessions Guide</i> .
<code>X-OpenAM-Username</code> , <code>X-OpenAM-Password</code>	Used to pass credentials in REST calls that use the HTTP POST method. See <i>"Authenticating (REST)"</i> in the <i>Authentication and Single Sign-On Guide</i> .
<code>Accept-API-Version</code>	Used to request a specific AM endpoint version. See <i>"REST API Versioning"</i> in the <i>Getting Started with REST</i> .
<code>Content-Type</code>	Required for cross-origin calls to AM REST API endpoints.
<code>If-Match</code> , <code>If-None-Match</code>	Used to ensure the correct version of a resource will be affected when making a REST call, for example when updating an UMA resource. See <i>"To Update an UMA Resource"</i> in the <i>User-Managed Access (UMA) 2.0 Guide</i> .

Exposed Headers

Add the response header names that AM returns in the `Access-Control-Expose-Headers` header.

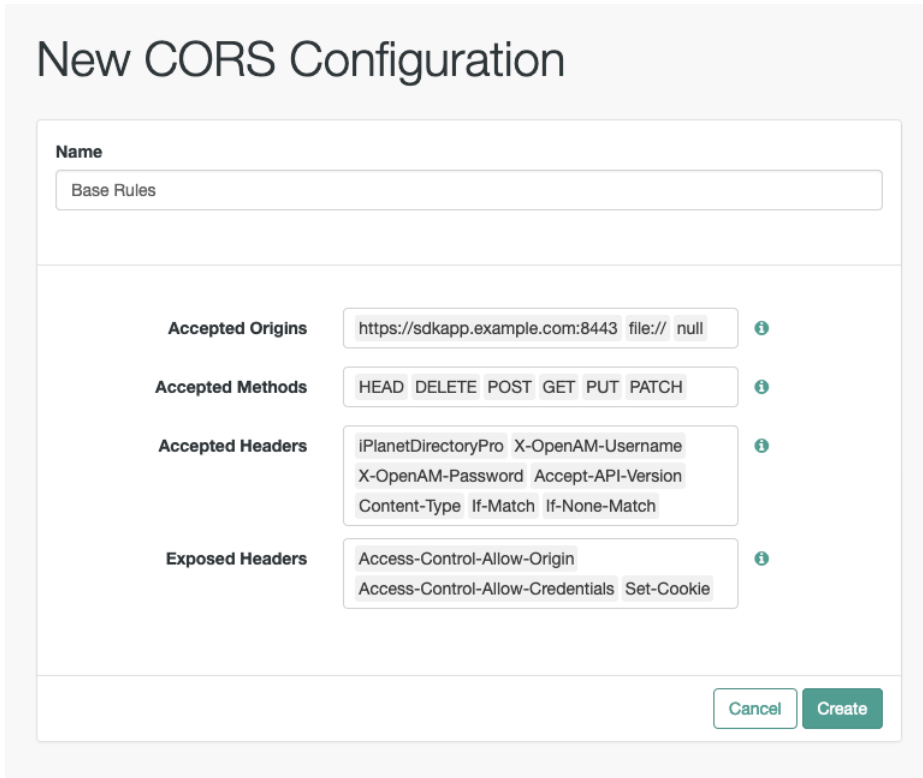
The header names are case-insensitive.

User agents can make use of any headers that are listed in this property, as well as the simple response headers, which are as follows:

- `Cache-Control`
- `Content-Language`
- `Expires`
- `Last-Modified`
- `Pragma`
- `Content-Type`

User agents must filter out all other response headers.

Example:



After you have completed the initial form fields, click Create.

The main CORS configuration page has the following additional properties:

Enable the CORS filter

Specifies whether the values specified in this CORS configuration instance will be active.

Max Age

The maximum length of time, in seconds, that the browser is allowed to cache the pre-flight response. The value is included in pre-flight responses, in the `Access-Control-Max-Age` header.

Allow Credentials

Whether to allow requests with credentials in either HTTP cookies or HTTP authentication information.

Enable this property if you send `Authorization` headers as part of the CORS requests, or need to include information in cookies when making requests.

When enabled, AM sets the `Access-Control-Allow-Credentials: true` header.

To Delete a CORS Configuration

To delete a CORS configuration, go to **Configure > Global Services > CORS Service > Secondary Configurations**, and next to the configuration to delete, click the **Delete (✕)** button.

Tip

You can disable a CORS configuration, and enable it again later, by selecting the rule and toggling the **Enable the CORS filter** property.

Configure CORS (REST)

You can use the endpoint to add multiple CORS configurations to AM, which are combined and used to ensure that only your trusted clients and applications can access your AM instance's resources.

For example, you could use the REST endpoint to add a base configuration, allowing a broad set of headers, and then add a stricter configuration; for example, for your OAuth 2.0 clients.

Tip

For information about the `/global-config/services/CorsService` endpoint, see the API Explorer available in the AM console.

See the following examples of managing CORS configuration in AM by using the REST endpoint:

- "To Add a CORS Configuration"
- "To Delete a CORS Configuration"

To Add a CORS Configuration

To *add* a new CORS configuration, create an HTTP PUT request to the `/global-config/services/CorsService` REST endpoint.

Note

You will require the SSO token of an administrative user; for example, `amAdmin`.

For information on obtaining an SSO token by using REST, see "Authenticating (REST)" in the *Authentication and Single Sign-On Guide*.

The body of the PUT call should contain the information that the CORS service uses. The available parameters for the request body are as follows:

enabled

Specifies whether the values specified in the CORS configuration instance will be active (**true**), or not (**false**).

Note

At least one instance must be enabled for AM to enforce CORS.

acceptedOrigins

A comma-separated list of the *origins* allowed when making CORS requests to AM. Wildcards are not supported; each value should be an exact match for the origin of the CORS request.

Example:

```
{
  "acceptedOrigins": [
    "http://example.com",
    "https://example.org:8433"
  ]
}
```

The CORS service automatically collects the values of the JavaScript Origins property in each OAuth 2.0 client configured, and adds them to an internal list of accepted origins. You do not need to add them manually.

Tip

During development you may not be using fully qualified domain names as the origin of a CORS request; for example, you are using the **file://** protocol locally.

If so, you can add these non-FQDN origins to the list; for example, **http://example.com**, **https://example.org:8433**, **file://**, **null**.

acceptedMethods

A list of HTTP methods allowed when making CORS requests to AM. The list is included in pre-flight responses, in the **Access-Control-Allow-Methods** header.

The method names are case-sensitive, ensure they are entered in all uppercase characters.

Example:

```
{
  "acceptedMethods": [
    "GET",
    "POST",
    "PUT",
    "PATCH",
    "OPTIONS",
    "DELETE"
  ]
}
```

acceptedHeaders

A list of request header names allowed when making CORS requests to AM. The list is included in pre-flight responses, in the `Access-Control-Allow-Headers` header.

The header names are case-insensitive.

Example:

```
{
  "acceptedHeaders": [
    "iPlanetDirectoryPro",
    "X-OpenAM-Username",
    "X-OpenAM-Password",
    "Accept-API-Version",
    "Content-Type",
    "If-Match",
    "If-None-Match"
  ]
}
```

By default, the following simple headers are explicitly accepted:

- `Cache-Control`
- `Content-Language`
- `Expires`
- `Last-Modified`
- `Pragma`

If you do not specify values for this element, the presence of any header in the CORS request, other than the simple headers listed above, will cause the request to be rejected.

+ *What are the commonly used headers?*

Headers commonly used when accessing an AM server include the following:

Commonly Used Headers

Header	Information
<code>iPlanetDirectoryPro</code>	Used for session information. See <i>"Introducing Sessions"</i> in the <i>Sessions Guide</i> .
<code>X-OpenAM-Username</code> , <code>X-OpenAM-Password</code>	Used to pass credentials in REST calls that use the HTTP POST method. See <i>"Authenticating (REST)"</i> in the <i>Authentication and Single Sign-On Guide</i> .
<code>Accept-API-Version</code>	Used to request a specific AM endpoint version.

Header	Information
	See " <i>REST API Versioning</i> " in the <i>Getting Started with REST</i> .
Content-Type	Required for cross-origin calls to AM REST API endpoints.
If-Match, If-None-Match	Used to ensure the correct version of a resource will be affected when making a REST call, for example when updating an UMA resource. See "To Update an UMA Resource" in the <i>User-Managed Access (UMA) 2.0 Guide</i> .

exposedHeaders

A list of response header names that AM returns in the `Access-Control-Expose-Headers` header.

The header names are case-insensitive.

User agents can make use of any headers that are listed in this property, as well as the simple response headers, which are as follows:

- Cache-Control
- Content-Language
- Expires
- Last-Modified
- Pragma
- Content-Type

User agents must filter out all other response headers.

Example:

```
{
  "exposedHeaders": [
    "Access-Control-Allow-Origin",
    "Access-Control-Allow-Credentials",
    "Set-Cookie"
  ]
}
```

maxAge

The maximum length of time, in seconds, that the browser is allowed to cache the pre-flight response. The value is included in pre-flight responses, in the `Access-Control-Max-Age` header.

allowCredentials

Whether to allow requests with credentials in either HTTP cookies or HTTP authentication information.

Set to `true` if you send `Authorization` headers as part of the CORS requests, or need to include information in cookies when making requests.

When enabled, AM sets the `Access-Control-Allow-Credentials: true` header.

The following shows an example of configuring CORS rules by using the `/global-config/services/CorsService` endpoint:

```
$ curl \
--request PUT \
--header "Content-Type: application/json" \
--header "X-Requested-With: XMLHttpRequest" \
--header "iplanetDirectoryPro: AQIC5w...2NzEz*" \
--data '{
  "enabled": true,
  "acceptedOrigins": [
    "http://localhost:8000",
    "null",
    "file://",
    "https://example.org:8443"
  ],
  "acceptedMethods": [
    "POST",
    "PUT",
    "OPTIONS"
  ],
  "acceptedHeaders": [
    "iPlanetDirectoryPro",
    "X-OpenAM-Username",
    "X-OpenAM-Password",
    "X-OpenIDM-Username",
    "X-OpenIDM-Password",
    "X-OpenIDM-NoSession",
    "Accept",
    "Accept-API-Version",
    "Authorization",
    "Cache-Control",
    "Content-Type",
    "If-Match",
    "If-None-Match",
    "X-Requested-With"
  ],
  "exposedHeaders": [
    "Access-Control-Allow-Origin",
    "Access-Control-Allow-Credentials",
    "WWW-Authenticate",
    "Set-Cookie"
  ],
  "maxAge": 1800,
  "allowCredentials": true
}' \
https://openam.example.com:8443/openam/json/global-config/services/CorsService

{
  "_id": "ef61e99c-6c83-4044-a1f5-71f472531b71",
  "_rev": "-1255664842",
  "maxAge": 1800,
  "exposedHeaders": [
```

```

    "Access-Control-Allow-Origin",
    "Access-Control-Allow-Credentials",
    "WWW-Authenticate",
    "Set-Cookie"
  ],
  "acceptedOrigins": [
    "null",
    "file://",
    "https://example.org:8443",
    "http://localhost:8000"
  ],
  "acceptedMethods": [
    "POST",
    "OPTIONS",
    "PUT"
  ],
  "acceptedHeaders": [
    "iPlanetDirectoryPro",
    "X-OpenAM-Username",
    "X-OpenAM-Password",
    "X-OpenIDM-Username",
    "X-OpenIDM-Password",
    "X-OpenIDM-NoSession",
    "Accept",
    "Accept-API-Version",
    "Authorization",
    "Cache-Control",
    "Content-Type",
    "If-Match",
    "If-None-Match",
    "X-Requested-With"
  ],
  "enabled": true,
  "allowCredentials": true,
  "_type": {
    "_id": "CORSService",
    "name": "CORS Service",
    "collection": true
  }
}

```

On success, AM returns an HTTP 201 response code, and a representation of the CORS settings, in JSON format. AM generates a UUID for the configuration, returned as the value of the `_id` property. You can use this ID value to update or delete the configuration with additional REST calls.

The new settings take effect immediately.

To Delete a CORS Configuration

To *delete* a CORS configuration, create an HTTP DELETE request to the `/global-config/services/CorsService` REST endpoint.

Note

You will require the SSO token of an administrative user; for example, `amAdmin`.

For information on obtaining an SSO token by using REST, see "Authenticating (REST)" in the *Authentication and Single Sign-On Guide*.

Add the ID of the configuration to delete to the URL.

The following shows an example of deleting CORS rules by using the `/global-config/services/CorsService` endpoint:

```
$ curl \
--request DELETE \
--header "X-Requested-With: XMLHttpRequest" \
--header "iplanetDirectoryPro: AQIC5w...2NzEz*" \
https://openam.example.com:8443/openam/json/global-config/services/CorsService/ef61e99c-6c83-4044-a1f5-71f472531b71

{
  "_id": "ef61e99c-6c83-4044-a1f5-71f472531b71",
  "_rev": "-1255664842",
  "maxAge": 1800,
  "exposedHeaders": [
    "Access-Control-Allow-Origin",
    "Access-Control-Allow-Credentials",
    "WWW-Authenticate",
    "Set-Cookie"
  ],
  "acceptedOrigins": [
    "null",
    "file://",
    "https://example.org:8443",
    "http://localhost:8000"
  ],
  "acceptedMethods": [
    "POST",
    "OPTIONS",
    "PUT"
  ],
  "acceptedHeaders": [
    "iPlanetDirectoryPro",
    "X-OpenAM-Username",
    "X-OpenAM-Password",
    "X-OpenIDM-Username",
    "X-OpenIDM-Password",
    "X-OpenIDM-NoSession",
    "Accept",
    "Accept-API-Version",
    "Authorization",
    "Cache-Control",
    "Content-Type",
    "If-Match",
    "If-None-Match",
    "X-Requested-With"
  ],
  "enabled": true,
  "allowCredentials": true,
  "_type": {
    "_id": "CORSService",
    "name": "CORS Service",
    "collection": true
  }
}
```

```
}  
}
```

On success, AM returns an HTTP 200 response code, and a representation of the CORS settings that were deleted, in JSON format.

The changes to the CORS settings take effect immediately.

Configuring the Cookie Domain

Configure AM's cookie domain to ensure only users and entities from trusted domains can be authenticated.

By default, the AM installer sets the cookie domain based on the fully qualified hostname of the server on which it installs AM, such as `openam.example.com`.

After installation, you may want to change the cookie domain to `example.com` so AM can communicate with any host in the sub-domain.

To Change the Cookie Domain

1. Log in to the AM console as an administrator user, for example, `amAdmin`.
2. Navigate to Configure > Global Services > Platform > Cookie Domain.
3. In the Cookie Domain field, set the list of domains into which AM should write cookies. Consider the following points:
 - Configure as many cookie domains as your environment requires. For example, for the realms configured with DNS aliases¹. Browsers ignore any cookies that do not match the current domain to ensure the correct one is used.
 - If you do not specify any cookie domain, AM uses the fully qualified name of the server, which implies that a host cookie is set rather than a domain cookie.

When configuring AM for Cross-Domain Single Sign-On (CDSSO), you must protect your AM deployment against cookie hijacking by setting a host cookie rather than a domain cookie. For more information, see "Enabling Restricted Tokens for CDSSO Session Cookies".

- Do not configure a top-level domain as your cookie domain; browsers will reject them. Top-level domains are browser-specific. For example, Firefox considers special domains like Amazon's web service (for example, `ap-southeast-2.compute.amazonaws.com`) to be a top-level domain².
- Do not configure the cookie domain such that it starts with a dot (.) character. For example, configure `example.com` instead of `.example.com`.

¹For more information, see "To Configure DNS Aliases for Accessing a Realm" in the *Setup Guide*.

²For a list of effective top-level domains, see https://publicsuffix.org/list/effective_tld_names.dat.

- If you are using Wildfly as the AM web container with multiple cookie domains, you must set the advanced server property, `com.sun.identity.authentication.setCookieToAllDomains`, to `false`.

Set this property in the AM console under Configure > Server Defaults > Advanced.

4. Save your changes.
5. Restart AM or the container where it runs.

Cross-Site Request Forgery (CSRF) Protection

AM includes a global filter to harden AM's protection against CSRF attacks. The filter applies to all REST endpoints under `json/` and requires that all requests other than GET, HEAD, or OPTIONS include, at least, one of the following headers:

- `X-Requested-With`

This header is often sent by Javascript frameworks, and the UI already sends it on all requests.

- `Accept-API-Version`

This header specifies which version of the REST API to use. Use this header in your requests to ensure future changes to the API do not affect your clients.

For more information about API versioning, see "*REST API Versioning*" in the *Getting Started with REST*.

Failure to include at least one of the headers causes the REST call to fail with a `403 Forbidden` error, even if the SSO token is valid.

Note

The CSRF filter applies *only* when the request includes the SSO token in the session cookie (`iPlanetDirectoryPro` by default).

To disable the filter, go to Configure > Global Services > REST APIs > and turn off Enable CSRF Protection.

The `json/` endpoint is not vulnerable to CSRF attacks when the filter is disabled, since it requires the "`Content-Type: application/json`" header, which currently triggers the same protection in browsers. This may change in the future, so it is recommended to enable the CSRF filter.

Chapter 3

Securing Administrative Access

Some deployments may require a single administrator, for example, those where the configuration never changes in production. However, if your deployment benefits from having more than one administrator user, it makes sense to limit what individual administrators can do.

This approach not only reduces the risk of accidental or intentional abuse of power, but also allows you to split the work between different teams and to audit changes in AM's configuration.

Securing administration in AM can be summed up as follows:

- Understanding the `amAdmin` user and learning how to delegate realm privileges to groups of users.
- Securing access to the AM administration console, and to the tools that you can use to configure AM: `Amster` and the `ssoadm` command.

About the amAdmin User

When you install AM, the `amAdmin` administrator account is created. This user is capable of unrestricted changes to AM configuration, including creating new users and augmenting their list of administrative privileges.

The `amAdmin` account cannot be deleted since it is hard-coded in the source code of several files. The user is defined within AM's configuration, so it is always available to AM even in the event that the identity stores become unavailable. Since it is not an identity defined in any identity store, it cannot use any capabilities that require a user profile, such as Device Match or push notifications.

The `com.sun.identity.authentication.super.user` advanced server property defines the DN of the `amAdmin` user. This property can be changed to a DN of a regular user that exists in any identity store configured in AM.

Changing the name of the `amAdmin` user might, however, affect the functionality of those files where the user name is hard-coded.

Secure the `amAdmin` user with a strong password and restrict its use as much as possible; delegate realm administration privileges to regular users instead.

Changing the amAdmin Password (Console)

In this section you will find procedures to change the password of the top-level administrator `amAdmin`, when:

- AM is configured using an external configuration store.

See "To Change the amadmin User's Password: External Configuration Store".

- AM is configured for evaluation and is using the embedded DS server as the configuration store.

See "To Change the amadmin User's Password: Embedded Configuration Store".

Tip

For a different way to change the **amAdmin** user's password, regardless of how the configuration store is configured, see "Changing the amAdmin Password (Secret Stores)".

To Change the amadmin User's Password: External Configuration Store

If AM is configured to use an external configuration store, perform the following steps to change the **amAdmin** user's password:

1. Log in to the AM console as the administrator, **amAdmin**.
2. Navigate to the user avatar (👤) in the top right hand corner of the UI.
3. Select Change Password.
4. Enter the current password in the Current password field.
5. Enter the new password in the New password and Confirm new password fields.
6. Save your changes.

If your deployment has multiple AM servers, the new password replicates across all servers.

To Change the amadmin User's Password: Embedded Configuration Store

1. Back up your deployment as described in "*Backing Up Configurations*" in the *Maintenance Guide*.
2. Log in to the AM console as the administrator, **amAdmin**.
3. Navigate to the user avatar (👤) in the top right hand corner of the UI.
4. Select Change Password.
5. Enter the current password in the Current password field.
6. Enter the new password in the New password and Confirm new password fields.
7. Save your changes.

When AM is configured to use the embedded DS server for the configuration store, you must change the passwords of the `uid=admin` user to match the new `amAdmin` password.

8. Change the `uid=admin` account's bind password in the AM configuration as follows:
 - a. Change the password for the configuration store binding:
 - i. Navigate to Deployment > Servers > *Server Name* > Directory Configuration.
 - ii. Enter the new bind password, which is the new `amAdmin` password, and save your changes.

Note

Changing the bind password of the configuration store updates the `configstorepwd` alias in the AM keystore file the next time AM starts.

- b. (Optional) If you use the embedded DS server as a data store, change the following bind passwords:

- i. Navigate to Realms > *Realm Name* > Identity Stores > embedded:
 - Enter the new bind password, which is the new `amAdmin` password, and save your changes.

Make this change in every AM realm that uses the embedded DS as an identity store.

- ii. Navigate to Realms > *Realm Name* > Services > Policy Configuration:
 - Enter the new bind password, which is the new `amAdmin` password, and save your changes.

Make this change in every AM realm that uses the embedded DS as a data store.

- iii. Navigate to Realms > *Realm Name* > Authentication > Modules, and select LDAP:
 - Enter the new bind password, which is the new `amAdmin` password, and save your changes.

Make this change in every AM realm that uses the embedded DS as a data store.

9. To change the `uid=admin` and the global administrator passwords in the embedded DS, see *Forgotten Superuser Password* in the *ForgeRock Directory Services Maintenance Guide*.

Changing the amAdmin Password (Secret Stores)

Another way to change the password of the `amAdmin` user is to use a special secret store.

This secret store does *not* show in the AM console, and the password of the `amAdmin` user stored in a secret *must* be salted and hashed. Encryption is optional, but *highly recommended*.

Important

If you supply the `amAdmin` password using secrets, you cannot change the password using the AM console unless you remove the secret configuration.

When you remove the secret configuration, the `amAdmin` password will revert to what it was before you configured the secrets.

You can provide the password of the `amAdmin` user in different secrets, as shown in the following procedure:

To Store the Password of the `amAdmin` User in a Secret

1. Salt and hash the new password of the `amAdmin` user.

You can use a script similar to the following one. Review the comments to understand the salt and hash requirements:

```
#!/usr/bin/env python3
import getpass
import os
import sys
import struct
import hashlib
import base64

if os.isatty(0):
    pwd = getpass.getpass()
    cnf = getpass.getpass('Confirm: ')
else:
    pwd = sys.stdin.buffer.readline().decode('utf-8').strip()
    cnf = pwd

if pwd != cnf:
    sys.exit("Password and confirmation don't match")

## Create some random bytes as the salt
salt = os.urandom(20)

## Hash the salt and the new password with a SHA-512 function
h = hashlib.sha512()
h.update(salt)
h.update(pwd.encode('utf-8'))
hash = h.digest()

## Concatenate the salt length as a single byte, the raw salt, and the hashed password
packed = struct.pack("B20s64s", 20, salt, hash)

## Generate the final hashed string
outform = "{SSHA-512}" + base64.b64encode(packed).decode('ascii')
print(outform)
```

2. Decide whether to encrypt the hashed string, and how to do it:

- Encrypting with the AM encryption password

1. Log in to the AM console with an administrative user. For example, `amAdmin`.
2. Go to <https://openam.example.com/openam/encode.jsp>, and paste the final hashed string in the field.

Optionally, you can use the **ampassword** command to encrypt the password. See "*Setting Up Administration Tools*" in the *Installation Guide* and `ampassword(1)` in the *Reference*.

3. Go to Configure > Server Defaults > Advanced.
4. Set the `org.forgerock.openam.secrets.special.user.passwords.format` advanced server property to `ENCRYPTED_PLAIN`.

- Encrypting with a secret stored in the Google Cloud KMS

+ *Prerequisites*

You need a Google Cloud Platform account that has a project. The project must have:

- A key ring containing the secrets that you will use to encrypt the hash of the password of the `amAdmin` user.

The key ring can be configured in any Google Cloud location.

- A service account that AM will use to connect to the project.

Refer to the [Google Key Management Service documentation](#) and Google's [Getting Started with Authentication](#) for more information.

To configure AM to connect to the Google Cloud KMS with the service account, see "[Configuring the Google Service Account Credentials](#)".

1. Check if you already have a Google Cloud KMS secret for decrypting.

Go to Configure > Server Defaults > Advanced, and check if the `org.forgerock.openam.secrets.googlekms.decryptionkey` advanced server property is configured.

If it is, you do not need to create another key.

If the property is not configured, log in to your Google Cloud dashboard and create a secret of one of the following types in the key ring of your choosing:

- Symmetric encrypt/decrypt
- Asymmetric decrypt

2. Use the secret you identified or created in the previous step to encrypt the hashed string.

You can use the **gcloud** tool included in Google Cloud's SDK to encrypt it. The tool creates a binary file with the encrypted secret, but AM does not support secrets in binary format. To work around this, base64-encode the encrypted secret. For example:

```
gcloud kms encrypt \
--plaintext-file=./amadmin_password_hashed_string.txt \
--ciphertext-file=- \
--project=my_project_ID \
--location=my_location \
--keyring=my_keyring_for_AM \
--key=my_key_for_decrypting_secrets_in_AM | base64 > encrypted_hash_of_amadmin_password.enc
```

3. In the AM console, go to Configure > Server Defaults > Advanced.
4. (Optional) If unset, set the `org.forgerock.openam.secrets.googlekms.decryptionkey` advanced server property to the fully qualified resource ID of the Google Cloud KMS secret that you used to encrypt the hash string. For example:

```
projects/my_project_ID/locations/my_location/keyRings/my_keyring_for_AM/cryptoKeys/
my_key_for_decrypting_secrets_in_AM
```

For information about how to find the key ID, see [Object Hierarchy](#) in the Google Cloud KMS documentation.

5. Set the `org.forgerock.openam.secrets.special.user.passwords.format` advanced server property to `GOOGLE_KMS_ENCRYPTED`.
- Leaving the hashed string unencrypted

Caution

Ensure that the password is randomly generated and high entropy before continuing.

1. In the AM console, go to Configure > Server Defaults > Advanced.
2. Set the `org.forgerock.openam.secrets.special.user.passwords.format` advanced server property to `PLAIN`.

Tip

If you cannot access the AM console, you can instead add the required property to the `CATALINA_OPTS` variable. For example, for Apache Tomcat, add the following to the `$CATALINA_BASE/bin/setenv.sh` file:

```
export CATALINA_OPTS="$CATALINA_OPTS -
Dorg.forgerock.openam.secrets.special.user.passwords.format=PLAIN"
```

3. Map the encrypted output to the secret ID that you will use. Perform one of the following:

- Save the encrypted password to a file in the special secret store directory:

```
$ echo -n amadmin_salted_encrypted_pass > /path/to/openam/security/secrets/userpasswords/
password.amadmin
```

Tip

The default location of the special secret store is `/path/to/openam/security/secrets/userpasswords`. To change it, configure the `org.forgerock.openam.secrets.special.user.passwords.dir` advanced server property.

- Store the encrypted password in an operating system variable called `PASSWORD_AMADMIN`, and make sure it is available to the user running the container where AM runs. For example, add it to the user's `bash.profile` file.
- Store the encrypted password in a Java system variable called `password.amadmin`, and make sure it is available to the container where AM runs.

For example, if using Apache Tomcat, add it to `$CATALINA_BASE/bin/setenv.sh` as follows:

```
export password.amadmin="y3GVzNP5Z3$EXZQHX75aRE!8FjN"
```

Delegating Privileges

The `amAdmin` user can change any setting in AM's configuration, but giving that power to each of your administrative users is not ideal.

In AM, you do not create administrative users. You create regular users and delegate realm administration privileges to a group they belong to. For example, you can create a group of users that are only allowed to make REST calls to endpoints to a specific realm, or a group of users that have full administration privileges for a particular realm.

This approach of splitting responsibilities lowers the risk of accidental or intentional abuse.

Since users with delegated administration privileges are regular users in the identity store, they can use any form of multi-factor authentication.

You can also delegate other kinds of privileges, such as making REST calls to realms for policy evaluation, modifying policies, and more.

+ *Realm Privileges Available for Delegation*

The following table describes privileges that you can assign in the AM console or by using the `ssoadm add-privileges` command:

Privileges

Privilege as it Appears in the AM console	Privilege Name to Use With the ssoadm add-privileges Command	Notes
Read and write access to all realm and policy properties	<code>Realm Admin</code>	Assign this privilege to administrators in order to let them modify or read any part of an AM realm. Use this privilege when you do not require granularity in your delegation model. All other AM privileges are included with this privilege. Administrators using the AM administration console must have this privilege.
Read and write access to all configured agents	<code>Agent Admin</code>	Provides access to centralized agent configuration; subset of the <code>RealmAdmin</code> privilege.
Read and write access to all log files	<code>Log Admin</code>	Subset of the <code>Realm Admin</code> privilege.
Read access to all log files	<code>Log Read</code>	Subset of the <code>Realm Admin</code> privilege.
Write access to all log files	<code>Log Write</code>	Subset of the <code>RealmAdmin</code> privilege.
Read and write access to all federation metadata configurations	<code>Federation Admin</code>	Subset of the <code>Realm Admin</code> privilege.
REST calls for reading realms	<code>Realm Read Access</code>	Subset of the <code>Realm Admin</code> privilege.
Read and write access only for policy properties, including REST calls	<code>Policy Admin</code>	Assign this privilege to policy administrators in order to let them modify or read any part of the AM policy configuration. This privilege lets an administrator modify or read all policy components: policies, applications, subject types, condition types, subject attributes, and decision combinators. All other AM privileges that affect policy components are included with this privilege. Subset of the <code>Realm Admin</code> privilege.
REST calls for policy evaluation	<code>Entitlement Rest Access</code>	Subset of the <code>Realm Admin</code> and <code>Policy Admin</code> privileges.
REST calls for reading policies	<code>Privilege Rest Read Access</code>	Subset of the <code>Realm Admin</code> and <code>Policy Admin</code> privileges.
REST calls for managing policies	<code>Privilege Rest Access</code>	Subset of the <code>Realm Admin</code> and <code>Policy Admin</code> privileges.

Privilege as it Appears in the AM console	Privilege Name to Use With the ssoadm add-privileges Command	Notes
REST calls for reading policy applications	Application Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for modifying policy applications	Application Modify Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading policy resource types	Resource Type Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for modifying policy resource types	Resource Type Modify Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading policy application types	Application Types Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading environment conditions	Condition Types Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading subject conditions	Subject Types Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading decision combiners	Decision Combiners Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for reading subject attributes	Subject Attributes Read Access	Subset of the Realm Admin and Policy Admin privileges.
REST calls for modifying session properties	Session Property Modify Access	Subset of the Realm Admin and Policy Admin privileges.


To Delegate Privileges

1. Navigate to the realm for which you want to delegate privileges. For example, navigate to Realms > Top Level Realm.

Important

Delegating *administrative* privileges in the Top Level Realm allows members of the group full access to the AM instance. Administration privileges in any other realm allows the group to access administrative functionality only in that realm, and any child realms.

2. Navigate to Identities > Groups and select the name of the group to which you intend to grant access. If you do not have a group yet, create one.

The  **All Authenticated Identities** virtual group allows you to assign privileges to any identity that has a valid session in AM. Use it with caution, because not every identity authenticates to AM by using strong authentication. For example, this virtual group may contain the anonymous user, which is enabled by default.

3. Select the administrative privileges to delegate for the realm:

- (Optional) To grant users in the group access to the administration console for the realm, select Realm Admin.

Administrators can use the AM administration console as follows:

- Delegated administrators with the **Realm Admin** privilege can access full administration console functionality within the realms they can administer.
- Users with lesser privileges, such as the **Policy Admin** privilege, can not access the AM administration console, but can use REST to create and manage the functionality for which they have privileges.
- Both the top level administrator (such as **amAdmin**) and delegated administrators in the Top Level Realm with the **Realm Admin** privilege have access to full console functionality in all realms and can access AM's global configuration.
- (Optional) To grant users in the group access to REST endpoints, select the required privileges from the list.

For information about the available AM privileges, see [Realm Privileges Available for Delegation](#).

4. Save your work.

To enable delegated subrealm administrators to invalidate sessions, you must add an attribute to their entry in the data store, as described in the following procedure:

To Enable Delegated Subrealm Administrators to Invalidate Sessions

1. Create an LDIF file that modifies the distinguished name entry of the subrealm administrator, adds the **iplanet-am-session-destroy-sessions** attribute, and sets its value to the subrealm's DN.

In the following example, the delegated administrator is named **subRealmAdmin** and the subrealm is called **mySubRealm**:

```
dn: uid=subrealmadmin,ou=people,dc=openam,dc=forgerock,dc=org
changetype: modify
add: objectClass
objectClass: iplanet-am-session-service
-
add: iplanet-am-session-destroy-sessions
iplanet-am-session-destroy-sessions: o=mysubrealm,ou=services,dc=openam,dc=forgerock,dc=org
```

Note

All values in the LDIF must be in lowercase, even if the subrealm or administrator name is not.

2. Run the **ldapmodify** command included with DS to apply the LDIF file to the user data store. For example:

```
$ /path/to/openssl/bin/ldapmodify \
--hostname 'id.example.com' \
--port 1636 \
--useSsl \
--usePkcs12TrustStore /path/to/openssl/config/keystore \
--trustStorePasswordFile /path/to/openssl/config/keystore.pin \
--bindDN uid=admin \
--bindPassword str0ngAdm1nPa55word \
/path/to/ldif.file
# Processing MODIFY request for uid=subrealmadmin,ou=people,dc=openam,dc=forgerock,dc=org
# MODIFY operation successful for DN uid=subrealmadmin,ou=people,dc=openam,dc=forgerock,dc=org
```

The delegated realm administrator will now be able to invalidate sessions created in the subrealm.

Delegating Agent Profile Creation

If you want to create agent profiles when installing web or Java agents, then you need the credentials of an AM user who can read and write agent profiles.

You can use the AM administrator account when creating agent profiles. If you delegate web or Java agent installation, then you might not want to share AM administrator credentials with everyone who installs agents.

To Create Agent Administrators for a Realm

Follow these steps to create *agent administrator* users for a realm:

1. In the AM console, browse to Realms > *Realm Name* > Identities.
2. On the Groups tab, click Add Group and create a group for agent administrators.
3. On the Privileges tab, select Realm Admin, and Save your changes.
4. Navigate to Realms > *Realm Name* > Identities. On the Identities tab, create as many agent administrator users as needed.
5. For each agent administrator user, edit the user profile.

On the Groups tab of the user profile, add the user to agent profile administrator group, and then Save your work.

6. Provide each system administrator who installs web or Java agents with their agent administrator credentials.

When installing Java agents with the `--custom-install` option, the system administrator can choose the option to create the profile during installation, and then provide the agent administrator user name and the path to a read-only file containing the agent administrator password. For silent installs, you can add the `--acceptLicense` option to auto-accept the software license agreement.

Securing the Administration Console

AM provides both end-user pages, located at `openam/XUI`, and an administrative console, located at `openam/ui-admin`.

Consider the following points to secure the AM console:

- Limit access to the administrative console.

For example, allow access to the console URI only to inbound connections from a specific network, or create a blacklist or a whitelist with the endpoints the console uses. For more information and some examples, see the [Knowledge Base](#).

- Ensure administrative users present strong-enough credentials when logging in to the AM administrative console.

By default, users that log to the console make use of the chain or tree configured in the Organization Authentication Configuration property for the realm. To locate this property, navigate to Realms > *Realm Name* > Authentication > Settings > Core.

Ensure that you change the default for all realms, including the Top Level Realm.

- Disable the AM API explorer in production environments.

To achieve this, navigate to Configure > Global Services > REST APIs, and selecting Disabled in the API Descriptors drop-down list.

Securing Administrative Tools

AM provides the following administrative tools that you can use instead of the administrative console to configure AM: **Amster** and **ssoadm**.

We recommend that you do not install the tools in the same server on same server as AM, so that administrators do not require a local system account on that server.

Also, we recommend that you create a username/password tree specifically for tools so you can track it easily in your logs.

Review the following information to secure access to the tools:

Amster

- If the administrative users connect to AM using interactive login, ensure that they present strong-enough credentials.

By default, users logging in through Amster make use of the chain or tree configured in the Administrator Authentication Configuration property for the realm. To locate this property, navigate to Realms > *Realm Name* > Authentication > Settings > Core.

- If the administrative users connect to AM using private key connections, make sure that you create your own keys and share them with AM. For more information, see the [Amster User Guide](#).

ssoadm

- By default, users logging in through the **ssoadm** command use of the chain or tree configured in the Administrator Authentication Configuration property for the realm. To locate this property, navigate to Realms > *Realm Name* > Authentication > Settings > Core.

Ensure that your administrative users present strong-enough credentials.

- The **ssoadm** command requires that you provide the password of the administrative user stored in cleartext in a file.

Ensure the file is read-only for its owner.

Chapter 4

Securing Realms

The AM installation process creates the Top Level Realm (/), which contains AM default configuration data. This realm cannot be deleted or renamed, since it is the root of the realm hierarchy in AM.

Consider the following list of security best practices related to realms:

Disable Module-based Authentication

Module-based authentication lets users authenticate using the `module=module-name` login parameter, therefore bypassing multi-factor authentication if multiple modules are configured in a chain with the same `authLevel`.

To disable module based authentication, navigate to Realms > *Realm Name* > Authentication > Settings > Security, and clear the Module Based Authentication check box.

Deactivate the Anonymous User

The anonymous user is enabled by default. To harden security, deactivate the anonymous user, unless anonymous access is specifically required in your deployment. See [How do I deactivate the default anonymous user in AM](#).

Create Strong Authentication Trees

Ensure your users log in to AM using sensible authentication trees, such as trees that enforce multi-factor authentication.

Configure Sensible Default Authentication Services

By default, users that log in to the console make use of the chain or tree configured in the Organization Authentication Configuration property for the realm. To locate this property, navigate to Realms > *Realm Name* > Authentication > Settings > Core.

Special care must be given when setting your *default* authentication tree or chain.

If you leave the default authentication to the `ldapService` chain, users can still post their username and password into the authentication endpoint to retrieve a session, regardless of the services configured for authentication.

For example, consider a deployment where you disable module-based authentication and keep the default authentication chain to the out-of-the-box `ldapStore` authentication chain using `DataStore` module. If you have set up two factor authentication for your users, your users can still access

their accounts without performing the correct two factor authentication chain login sequence by using the default `ldapService` chain.

When you set the default authentication tree or chain, make sure it is set to your most secure tree or chain once you are ready to go to production and not left to the default `ldapService` chain.

Ensure that you change the default for all realms, including the Top Level Realm.

About the Demonstration User

When installing AM for evaluation, using the embedded DS server, a `demo` user is created. This is a regular account with no administrative permissions and is intended for test and demo purposes. You should remove it from production environments.

To remove the `demo` account, navigate to Realms > Top Level Realm > Identities, select the `demo` account, and select `Delete`.

Chapter 5

Configuring Secrets, Certificates, and Keys

Encryption makes it possible to protect sensitive data, encoding it in such a way that only authorized parties can access it.

Signing allows the receiver of a piece of data to validate the sender's identity and ensures that the data has not been tampered with.

AM depends on signing and encryption to protect network communication and to keep data confidential and unalterable. In turn, signing and encryption depend on keys or secrets, which are generated using cryptographic algorithms.

AM uses the following methods to store keys or secrets:

- **The AM keystore.** Used by some features, it can be configured globally so its configuration is shared by any AM instance in a deployment, or individually per server.

AM also uses this keystore to start up.

During installation, AM deploys a JCEKS and a JKS keystore with several self-signed key aliases, for demo and test purposes only.

- **AM secret stores.** Introduced in 6.5, secret stores are repositories for cryptographic keys and credentials. They can be configured globally, so the configuration is shared by any instance on the site, or by realm.

Note

AM is migrating features from using the AM keystore to use secret stores.

AM supports configuring JVM system properties, key aliases stored in keystores or HSM, or files stored in filesystems or secret volumes, as secrets.

About the Default Keystores and Secret Stores

During installation, AM deploys a JKS and a JCEKS keystore with several self-signed key aliases for demo and test purposes only.

Both the AM keystore and the default secret stores use the default JCEKS keystore. The JKS keystore is not used by default, and can be safely deleted.

Do not use the default keys, keystores, and secret stores in production environments.

+ About the Default JCEKS and JKS Keystore Keys and Aliases

	JCEKS	JKS
Used by default in AM?	Yes ^a	No
In which path is it?	/path/to/openam/security/keystores/keystore.jceks	/path/to/openam/security/keystores/keystore.jks
Where is its password stored? ^b	/path/to/openam/security/secrets/default/.storepass	/path/to/openam/security/secrets/default/.storepass
Which test aliases does it contain?	es256test ^c es384test ^c es512test ^c hmacsigningtest ^d directenctest ^e rsajwtsigningkey ^f selfserviceenctest ^f selfservicesigntest ^g test ^f	test ^f
Which password strings does it contain?	configstorepwd ^h dsameuserpwd ⁱ	None
Where is the private key password file? ^j	/path/to/openam/security/secrets/default/.keypass	/path/to/openam/security/secrets/default/.keypass

^a New AM installations use the JCEKS keystore as the default keystore.

^b The password of the JCEKS and JKS keystores is a random-generated string stored in cleartext.

^c ECDSA key.

^d Symmetric HMAC key.

^e Symmetric Direct AES encryption key.

^f Asymmetric RSA key.

^g Symmetric secret signing key.

^h The value of the `configstorepwd` is a string. It is the password of the configuration store, which is accessed during AM startup.

ⁱ The value of the `dsameuserpwd` is a string. It is the password of a reserved service account, which is accessed during AM startup.

^j The password for all the key aliases in the JCEKS and JKS keystores is `changeit`, stored in cleartext.

+ About the Default Secret Stores

- `default-keystore`. This keystore-type secret store is mapped to the default JCEKS keystore.

It also contains ID mappings for several of the AM features that use keys.

For more information about the mappings, see "Secret ID Default Mappings".

- **default-password-store**. This filesystem-type secret store is mapped to `/path/to/openam/security/secrets/encrypted` and it is used to provide the passwords to open the **default-keystore** secret store:
 - The **storepass** file contains the encrypted password of the keystore.
 - The **entrypass** file contains the encrypted password of the keys inside the keystore.

Note

This password configuration is very similar to the one for the default JCEKS keystore. However, the files are different. While the password files for the JCEKS are in cleartext, the password files for the default secret store are encrypted with AM's encryption key.

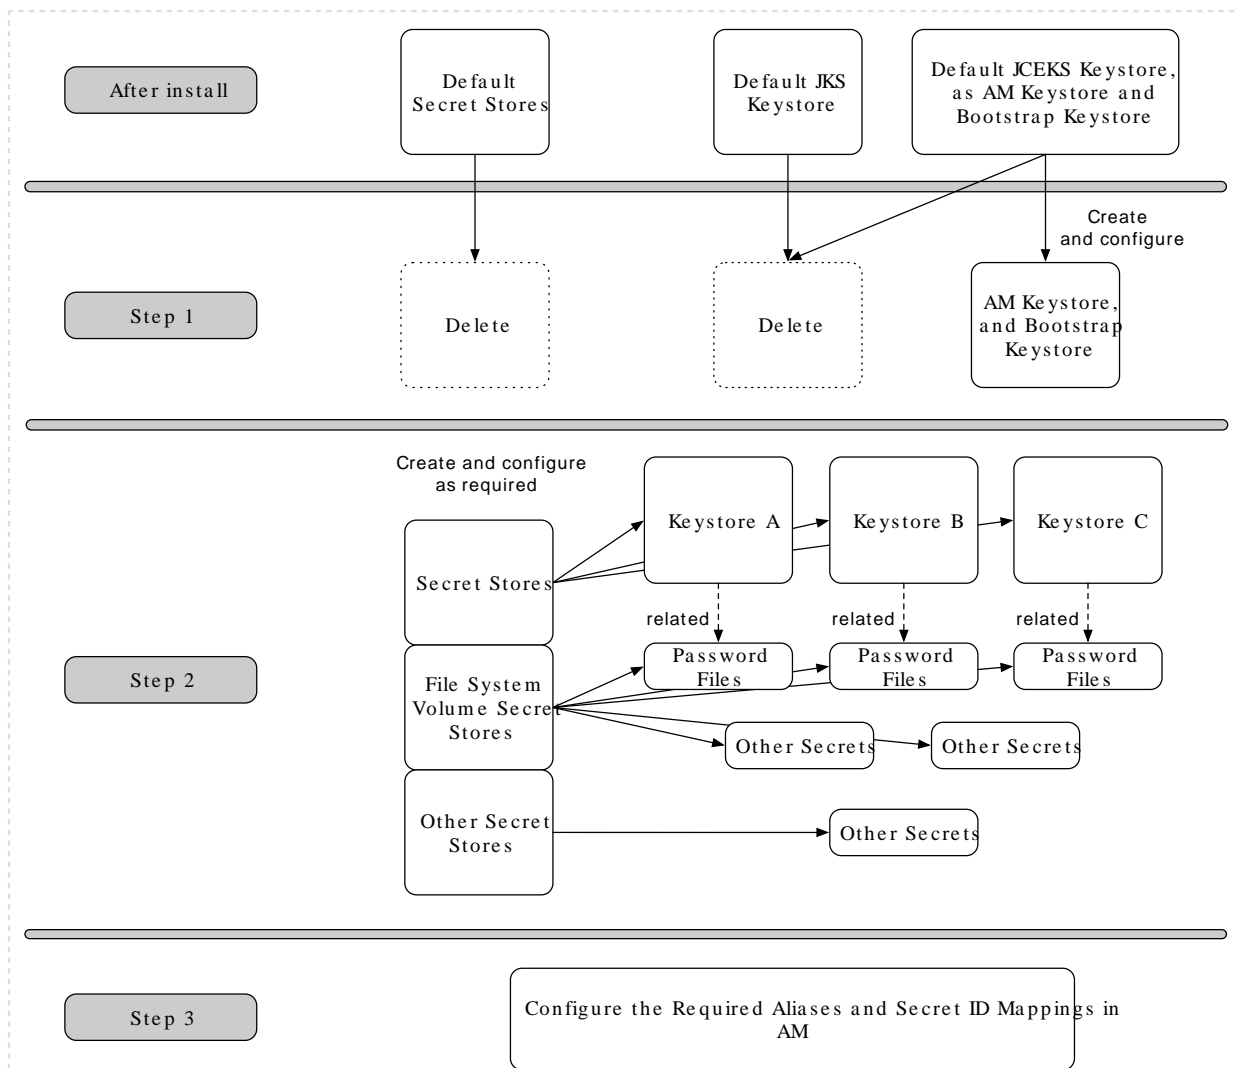
Take into account that the keystore file is the same, which means that if you change the passwords for the JCEKS keystore, you must change them in the default secret store as well.

Tasks to Configure Keystores and Secret Stores

The provided keystores and secret stores are sufficient for testing and demonstrating AM features.

For production and pre-production environments, configure the keystores and secret stores that your environment will use *before* configuring the AM features that use them.

+ *High-Level Steps to Configure Keystores and Secret Stores in Production Environments*



1. Create a new keystore to configure as the AM keystore, and delete the default keystores and the default secret stores.

AM will use the new keystore to start up.

Configure different keys for different features, when possible.

2. Create separate secret stores for each AM feature you are using. For example, create one for SAML v2.0 secrets, and a different one for OAuth 2.0 secrets.

Use different passwords among keystores or secret stores. This will reduce the amount of compromised keys or secrets if a malicious user is able to leverage one of the passwords.

Keystore secret stores need, at least, another secret store to provide the password of the keystore, and the password for the keys. For example, a file system volume secret store.

3. Configure AM features to use your custom key aliases and secrets.

The following table guides you through the tasks you need to perform to configure the keys and secrets AM requires:

Task	Resources
<p>Understand AM's Secret Needs</p> <p>Review the list of features that use keys in AM, and their possible keystore and secrets configurations.</p>	<ul style="list-style-type: none"> "Features in AM That Use Keys"
<p>Create a new AM Keystore and Configure It</p> <p>Create and configure a new AM keystore, which will also serve as the AM bootstrap keystore, and delete the default keystores and secret stores.</p>	<ul style="list-style-type: none"> "Managing the AM Keystore"
<p>Create Secrets as Needed</p> <p>Create as many keystores, key aliases, and/or secrets as required in your environment based on the information you learned when you reviewed the list on the first task. You will configure them in AM in the next steps.</p> <p>Keys and secrets protect the credentials, tokens, and other sensitive information that your environment needs to send and receive. Therefore, ensure that keys and secrets are protected and only shared when required. This may result in configuring multiple keystores and/or secret stores for different features.</p> <p>Do not reuse passwords among keystores or secret stores. This will reduce the amount of compromised keys or secrets if a malicious user is able to leverage one of the passwords.</p>	<ul style="list-style-type: none"> "Managing Key Aliases and Passwords"
<p>Configure Secret Stores in AM</p> <p>Create new secret stores to map the relevant new keys you created in previous tasks, for example, those for the OAuth 2.0 providers.</p>	<ul style="list-style-type: none"> "Configuring Secret Stores"
<p>Make Available the Keystores and Secret Stores to All AM Instances</p>	<p>Keystores and secret stores must be available in the same location across of all of the instances.</p>

Task	Resources
	This step may mean mounting a filesystem with the required files across the instances, installing cryptographic cards, and others, if not done already.
Configure Key Aliases and Secrets in AM Change the pre-configured key aliases and secrets with those created in the previous tasks.	<ul style="list-style-type: none"> • "Changing Default Key Aliases" • "Mapping and Rotating Secrets"

Features in AM That Use Keys

Most features that require storing secrets for signing or encryption use the AM keystore, which is configured by going to [Configure > Server Defaults > Security > Key Store](#). However, some features may require or support different configurations:

Features that only use the AM keystore

- **User self-service**

Requires a JCEKS keystore with a key pair alias for encryption and a key alias for signing. For more information, see ["Creating a User Self-Service Service Instance"](#) in the *User Self-Service Guide*.

- **Client-based sessions**

Requires an key pair alias for encryption and a key pair alias for signing. For more information, see ["Configuring Client-Based Session Security"](#).

- **Persistent Cookie nodes (authentication trees)**

Requires a key pair alias for encryption. For more information, see ["Set Persistent Cookie Node"](#) in the *Authentication and Single Sign-On Guide*.

- **Amster**

Requires a `sms.transport.key` key alias to export and import encrypted passwords. For more information, see the [Amster Command-line Interface Guide](#).

- **IDM user self-registration**

Requires copying signing and encryption keys from the IDM installation into the AM keystore. For more information, see ["To Delegate User Self-Registration to IDM"](#) in the *User Self-Service Guide*.

Features that use secret stores

- **Web and Java agents**

Web Agents and Java Agents communicate with AM using a built-in OAuth 2.0 provider configured globally in AM. This communication requires a key alias for signing tokens. For more information, see the *ForgeRock Web Agents User Guide* and the *ForgeRock Java Agents User Guide*.

- **OAuth 2.0 providers**

Requires a key alias for signing client-based tokens and OpenID Connect ID tokens. For more information, see "Configuring Client-Based OAuth 2.0 Token Digital Signatures" in the *OAuth 2.0 Guide* and "Configuring ID Token Signatures" in the *OpenID Connect 1.0 Guide*.

Also requires a key alias for direct authentication encryption of client-based OAuth 2.0 access and refresh tokens. For more information on enabling encryption, see "Configuring Client-Based OAuth 2.0 Token Encryption" in the *OAuth 2.0 Guide*.

- **Persistent cookie modules (authentication chains)**

Requires a key pair alias for encryption. For more information, see "Persistent Cookie Module" in the *Authentication and Single Sign-On Guide*.

- **Remote Consent Service**

Requires a key alias for signing consent responses, and another key alias for encrypting consent responses. For more information, see "*The Remote Consent Service*" in the *OAuth 2.0 Guide*.

- **Authentication trees**

Requires a key alias to encrypt values stored in the authentication tree's secure state. For more information, see "Storing Secret Values in Transient Tree State and the Secure State" in the *Authentication Node Development Guide*.

- **SAML v2.0 Federation**

Requires key pairs for signing and encryption of messages, responses, and assertions; for example, a key to encrypt the JWT stored in the local storage of supported browsers.

You may also require a key to sign exported meta data.

For more information, see "*Signing and Encryption*" in the *SAML v2.0 Guide*.

For a list of the secret ID mappings, see "Secret ID Default Mappings".

Features that support different keystore configurations:

- **ForgeRock Authenticator (OATH), ForgeRock Authenticator (PUSH) modules, and the WebAuthn Profile Encryption Service**

Supports configuring a different keystore to encrypt device profiles. They also support different keystore types that are not available to other features. For more information, see "About Multi-Factor Authentication" in the *Authentication and Single Sign-On Guide*.

- **AM's startup (bootstrap) process**

Requires two password strings. ForgeRock recommends that you use the AM keystore as the bootstrap keystore, but you can configure a bootstrap keystore as long as:

- You keep the password strings updated.
- You overwrite the `boot.json` file before AM starts up.

For more information, see "To Replace the AM Keystore".

Features that require different keystore configurations:

- **Java Fedlets**

Require a keystore containing a key pair to sign and verify XML assertions and to encrypt and decrypt SAML assertions. Keystore and key information are configurable in the `FederationConfig.properties` file. For more information, see "Configuring Java Fedlet Properties" in the *SAML v2.0 Guide*.

- **Security Token Service**

Requires configuring a JKS keystore for encrypting SAML v2.0 and OpenID Connect tokens. It does not require files to store the keystore password or the key aliases' passwords. For more information, see "*Configuring STS Instances*" in the *Security Token Service (STS) Guide*.

- **CSV audit logging handler**

Requires configuring a keystore for tamper-proofing. It does not require a file to store the keystore password; the password is configured in the AM console. For more information, see "Configuring CSV Audit Event Handlers".

Tip

If you are creating your own custom components or plugins, you can implement the `SecretIdProvider` interface for exposing your own custom secrets.

For more information, see the AM 7.0.2 Public API Javadoc.

Managing the AM Keystore

By default, AM installations provide a JCEKS keystore containing several test-only key aliases that are preconfigured in AM.

For production deployments, generate a new keystore with the key aliases you need to use.

+ *Considerations Before You Start*

- Different AM features support different keystore configurations, and some features do not use the default keystore to store their key aliases. For more information, see "Features in AM That Use Keys".
- Key aliases are not migrated from one keystore to another when changing the keystore configuration in AM. You must prepare the new keystore as required before configuring it.
- You must restart AM if you make any changes to the keystore. Changes are, for example, adding or removing keys or changing key or keystore passwords.

Tasks:

- To create a new AM keystore, see "To Replace the AM Keystore".
- To modify the AM keystore configuration without changing the content of the file, see "To Modify the AM Keystore Properties".

To Replace the AM Keystore

The AM keystore provides secrets to several features of AM, as explained in "Features in AM That Use Keys", but it also lets AM start up.

+ *The AM Keystore as the Bootstrap Keystore*

The AM startup process checks a file, called `boot.json`, that contains the settings AM requires to bootstrap. Among these settings are the path to a keystore file and the files containing the keystore and key passwords.

During the startup process, AM needs to find the following aliases inside the keystore configured in the `/path/to/openam/config/boot.json` file:

+ *About the `dsameuserpwd` and the `configstorepwd` Aliases*

configstorepwd

An alias for the password of the AM configuration store. The alias is password-protected, with the password specified by default in the `/path/to/openam/security/secrets/default/.keypass` file.

To update the value of this alias, go to Deployment > Servers > Server Name > Directory Configuration, and modify the configuration store bind password. Every time you change the bind alias, AM modifies the content of the key alias in the keystore file.

dsameuserpwd

An alias for the password of a special user required at AM startup time. The alias is password-protected, with the password specified by default in the `/path/to/openam/security/secrets/default/.keypass` file.

These strings cannot be recreated manually, but AM will recreate them in a new keystore after a successful start.

The key to successfully changing the AM/bootstrap keystore is to configure the new AM keystore, and restart AM while the old keystore is still accessible. The boot process will use the original keystore for booting up, write the password strings in the new keystore, and rewrite the `boot.json` file.

Follow the steps in this procedure to create a new AM/bootstrap keystore.

Perform the steps in this procedure to create a new keystore containing the password strings that AM needs to start up, and configure it as the new AM keystore:

1. Ensure that AM is running, and that you can access its console as an administrative user.
2. Acquire a new key from your certificate authority and add it to a new keystore, or generate a new self-signed key in a new keystore.

Create the keystore in a directory of your choosing. This directory should be the same for all the instances on the site. For example, `/path/to/openam/security/keystores`.

This example creates a self-signed key alias in a new keystore file, `am_keystore.jceks`, with a new asymmetric RSA key alias, `newkey`.

Note that in production environments, you should use the strongest algorithm you can.

```
$ keytool \
-genkeypair \
-alias newkey \
-keyalg RSA \
-keysize 2048 \
-validity 730 \
-storetype JCEKS \
-dname 'CN=newkey' \
-keystore am_keystore.jceks
Enter keystore password:
Reenter new password:
Enter key password for <newkey>
(RETURN if same as keystore password):
Reenter new password:
```

Take note of the passwords you entered.

3. Store the keystore passwords in *cleartext* in a directory of your choosing.

This directory should be the same for all the instances on the site. For example, `/path/to/openam/security/secrets/default`.

For example:

```
$ echo -n newstorepassword > .am_keystore_storepass
$ echo -n newkeypassword > .am_keystore_keypass
```

Use **echo -n** to avoid inserting hidden trailing newline characters. Even if the **keytool** command is able to use the password in the file, AM may not be able to open the keystore or the key aliases.

4. Make sure the password files have read-only permission for their owner. For example:

```
$ chmod 400 .am_keystore_storepass
$ chmod 400 .am_keystore_keypass
```

5. Configure the new keystore as the AM keystore in the site. Follow the steps in "To Modify the AM Keystore Properties".

Once AM starts, the new keystore contains the password strings that AM uses to start up. You can delete the default JCEKS keystore now. The default secret stores also use the JCEKS keystore. You can also delete them now.

To Modify the AM Keystore Properties

To modify the AM keystore configuration, perform the following steps:

1. In the AM console, go to Configure > Server Defaults > Security > Key Store.
2. Enter the keystore file name and path in the Keystore File field. For example, `/path/to/openam/security/kestores/am_keystore.jceks`.
3. Enter the Keystore Type. For example **JKS**, **JCEKS**, **PKCS11**, or **PKCS12**.
4. In the Keystore Password File field, enter the location of the keystore password file. For example, `/path/to/openam/security/secrets/default/.am_keystore_storepass`.
5. In the Private Key Password File field, enter the location of the private key password file. For example, `/path/to/openam/security/secrets/default/.am_keystore_keypass`.
6. Save your changes.

At this point, AM still holds the old keystore configuration in memory, and cannot use key aliases contained in the new keystore.

7. (Optional) If you need to change key aliases in the AM configuration, decide whether to change them before or after restarting the AM instances in the next step.

If you are using client-based sessions, ensure the signing key exists in the new keystore. You can check the configuration for client-based sessions by going to Realms > *Realm Name* > Authentication > Settings > Security.

To configure the rest of the features that need key aliases before or after the restart, see "Changing Default Key Aliases".

8. Make the new keystore files available in the same location to all the instances in the site.

This step may mean mounting a filesystem with the required files across the instances, copying the files across instances, and others.

9. Restart the AM instance or instances.

The new default keystore and its keys are ready to use.

Managing Key Aliases and Passwords

Whether you need to create new key aliases because you are using more AM features or because you are installing a new environment, consider the following points:

- Review the list of features in AM to understand which features use the AM keystore and which ones do not, then manage your secrets accordingly. You should avoid sharing certificates among features when possible, which may result in configuring different keystores or secret stores.

For more information, see "Features in AM That Use Keys".

- Make sure keystores, key aliases, and certificates are maintained on every instance; in a site environment, every instance has its own keystore files.
- Make sure keystores and secret stores are in the same location across of all of the instances on the site.

The following table contains a list of the tasks you may need to perform when managing key aliases in your environment:

Task	Resources
Create new key alias in an existing keystore or in a new keystore	"Creating Key Aliases"
Copy key aliases between keystores; for example, when configuring IDM's provisioning.	"Copying Key Aliases"
Change key alias passwords.	"Changing Key Alias Passwords"
Change keystore passwords.	"Changing Keystore Passwords".

Creating Key Aliases

Several AM features require key aliases for signing and encryption. New AM installations pre-configure default key aliases for all of the features, but you should create new key aliases for each of them.

You can create key aliases in a new keystore that, for example, will be configured later as the AM keystore, or you can create key aliases in the existing AM keystore:

- "To Create a Keystore and Key Aliases for Keystore Secret Stores"
- "To Create Key Aliases in an Existing Keystore"
- "To Create Self-Service Key Aliases"

To Create a Keystore and Key Aliases for Keystore Secret Stores

To create a new AM keystore, see "Managing the AM Keystore" instead.

1. Acquire a new key from your certificate authority and add it to a new keystore, or generate a new self-signed key in a new keystore.

This example creates a self-signed key alias in a new keystore file, `keystoreA.jceks`, with a new asymmetric RSA key alias, `newkey`.

Note that in production environments you should use the strongest algorithm you can.

```
$ cd /path/to/openam/security/keystores/
$ keytool \
-genkeypair \
-alias newkey \
-keyalg RSA \
-keysize 2048 \
-validity 730 \
-storetype JCEKS \
-dname 'CN=newkey' \
-keystore keystoreA.jceks
Enter keystore password:
Reenter new password:
Enter key password for <newkey>
(RETURN if same as keystore password):
Reenter new password:
```

Take note of the passwords. You need to make them available within another secret store; for example, by using a file system volume secret store, as shown below:

- a. Go to the directory that the filesystem volume secret store will point to. For example, `/path/to/openam/security/secrets/mydir`.

Since you can encode the content of the files using different modes, we recommend that you create a directory for each encode mode you plan to use, at least.

- b. Create two files, one for the keystore password, and another for the password of the keys inside the keystore.

The files will contain the passwords encoded expected by the file system secret volume store.

For example, if you choose `Base64 encoded` as the encoding, you must base64-encode the passwords, and then add them to their respective files.

For example:

```
$ echo -n bmV3c3RvcnVwYXNzd29yZA== > keystoreA_storepass
$ echo -n bmV3a2V5cGFzc3dvcmQ= > keystoreA_keypass
```

Important

Use **echo -n** to avoid inserting hidden trailing newline characters. Even if the **keytool** command is able to use the password in the file, AM may not be able to open the keystore or the key aliases.

- c. Make sure the password files have read-only permission for their owner. For example:

```
$ chmod 400 keystoreA_storepass
$ chmod 400 keystoreA_keypass
```

2. Create any other keys and keystores required by your environment by repeating the steps in this procedure and/or following the steps in "To Create Key Aliases in an Existing Keystore".
3. Ensure that password files and keystores are maintained on every instance in your environment. Every AM instance has its own keystores and password files.
4. Configure the keystore in a keystore-type secret store. See "Keystore Secret Stores".

To configure the file system secret store too, see "File System Secret Volumes Secret Stores".

To Create Key Aliases in an Existing Keystore

Perform the following steps to create new key aliases in an existing keystore. For example, the AM keystore:

1. Change directories to the keystore location, for example, `/path/to/openam/security/keystores/`.
2. Acquire a new key from your certificate authority, or generate a new self-signed key.

When you create or import a new key, the **keytool** command adds the new alias to the specified keystore if it exists, or creates a new keystore if it does not exist.

This example creates a self-signed key alias in the AM keystore, `am_keystore.jceks`, with a new asymmetric RSA key alias called `mynewkey`.

Note that in production environments you should use the strongest algorithm you can use.

```
$ cd /path/to/openam/security/keystores/
$ keytool \
-genkeypair \
-alias mynewkey \
-keyalg RSA \
-keysize 2048 \
-validity 730 \
-storetype JCEKS \
-dname 'CN=mynewkey' \
-keystore am_keystore.jceks
Enter keystore password: Enter the password in the .keystore_storepass file.
Enter key password for <mynewkey>
(RETURN if same as keystore password): Enter the password in the .keystore_keypass file.
Reenter new password: Enter the password in the .keystore_keypass file.
```

Remember:

- The contents of the password files of the AM keystore are in cleartext.
 - The contents of the password files in a file system volume secret store are *not* in cleartext by default. This means that you need to decode them before you can use them in the **keytool** command.
3. Ensure that password files and keystores are maintained on every instance in your environment. Every AM instance has its own keystores and password files.
 4. (AM keystore) Restart the AM instances affected by the configuration changes to use the new key aliases.
 5. Configure the new key aliases in AM. For a list of features that use key aliases and links to their relevant sections, see "Features in AM That Use Keys".

To Create Self-Service Key Aliases

User self-service requires a key pair for encryption and a signing secret key to be available in the AM keystore before configuring any of its features. Follow the steps in this procedure to create new key aliases for the user self-service features in the AM keystore:

1. Acquire a new key from your certificate authority, or generate new self-signed keys. The password of the new keys for the user self-service features must match the passwords of those keys already present in the keystore, and configured in the `/path/to/openam/security/secrets/default/.am_keystore_keypass` file.

This example generates a self-signed key for encryption and a new signing secret key in the `am_keystore.jceks` keystore, but you could also import CA-provided keys to the keystore.

- a. Create the new self-signed encryption key alias:

```
$ cd /path/to/openam/security/keystores/
$ keytool \
-genkeypair \
-alias newenckey \
-keyalg RSA \
-keysize 2048 \
-validity 730 \
-storetype JCEKS \
-dname 'CN=newenckey' \
-keystore am_keystore.jceks
Enter keystore password: Enter the password in the .am_keystore_storepass file.
Enter key password for <newenckey>
(RETURN if same as keystore password): Enter the password in the .am_keystore_keypass file.
Reenter new password: Enter the password in the .am_keystore_keypass file.
```

- b. Create the new signing secret key alias:

```
$ cd /path/to/openam/security/keystores/
$ keytool \
-genseckey \
-alias newsigkey \
-keyalg HmacSHA256 \
-keysize 256 \
-storetype JCEKS \
-keystore am_keystore.jceks
Enter keystore password: Enter the password in the .am_keystore_storepass file.
Enter key password for <newsigkey>
(RETURN if same as keystore password): Enter the password in the .am_keystore_keypass file.
Reenter new password: Enter the password in the .am_keystore_keypass file.
```

2. Ensure that password files and keystores are maintained on every instance in your environment. Every AM instance has its own keystores and password files.
3. Restart the AM instances affected by the configuration changes.
4. Configure user self-service to use the new keys. For more information, see "Creating a User Self-Service Service Instance" in the *User Self-Service Guide*.

Copying Key Aliases

Some AM features require access to the key aliases used by other components of the ForgeRock Identity Platform. For example, the IDM Provisioning feature requires access to the key aliases IDM uses for signing and encrypting data.

This section covers copying key aliases from the keystore of a ForgeRock Identity Platform component to AM's default keystore.

To Copy a Key Alias From One Keystore to Another

Use the **keytool** command to export the key from the source component's keystore. Install the result in the AM keystore, by performing the following steps:

1. From the source keystore, export the required key into a temporary keystore that can be transported to AM by executing the following **keytool** command:

```
$ keytool -importkeystore -srcstoretype jceks -srcalias "myKeyAlias" \
-deststoretype jceks -destalias "myKeyAlias" \
-srckeystore "/path/to/openidm/security/keystore.jceks" \
-destkeystore "/path/to/openidm/security/temp_keystore.jceks" \
-srckeypass "changeit" \
-srcstorepass "changeit" \
-destkeypass "myT3mPK3yP4ssword" \
-deststorepass "myT3mPK3yP4ssword"
```

This command exports the **myKeyAlias** key alias, specified by the **srcalias** argument, to a temporary keystore file **/path/to/openidm/security/temp_keystore.jceks**. The store and key password is set to **myT3mPK3yP4ssword**. You need to use the temporary passwords when importing to the AM instance.

2. Move the temporary keystore file created in the previous step, in this example **temp_keystore.jceks**, to the filesystem of the target AM server.
3. On the target AM server, import the key alias into the AM keystore by executing the following **keytool** command:

```
$ keytool -importkeystore -srcstoretype jceks -srcalias "myKeyAlias" \
-deststoretype jceks -destalias "myKeyAlias" \
-srckeystore "/path/to/openam/security/keystores/temp_keystore.jceks" \
-destkeystore "/path/to/openam/security/keystores/am_keystore.jceks" \
-srckeypass "myT3mPK3yP4ssword" \
-srcstorepass "myT3mPK3yP4ssword" \
-destkeypass:file "/path/to/openam/security/secrets/default/.am_keystore_keypass" \
-deststorepass:file "/path/to/openam/security/secrets/default/.am_keystore_storepass"
```

This command imports the key alias from the temporary **temp_keystore.jceks** keystore file, which was copied from the IDM instance, into the AM keystore. The command also sets the passwords to match those used by the default AM keystore.

4. (Optional) Repeat the previous steps to copy any additional key aliases from the source keystore to the destination keystore.
5. Restart the AM instance for the key change to take effect.

The AM instance will now be able to correctly encrypt, decrypt, sign or verify data and share it with the source ForgeRock Identity Platform component.

Changing Key Alias Passwords

Decrypting a key alias in a keystore requires a password. This password is initially specified when you generate the key, or when you import the key into a keystore, but you might need to update the password at a later time.

To Change Key Alias Passwords

1. Back up your keystore and password files.

2. (AM keystore) Replace the old password in the `.am_keystore_keypass` file with the new one:

```
$ echo -n newpassword > /path/to/openam/security/secrets/default/.am_keystore_keypass
```

Important

Use `echo -n` to avoid inserting hidden trailing newline characters. Even if the `keytool` command is able to use the password in the file, AM may not be able to use the key aliases if there are hidden trailing newline characters in the password file.

3. (Secret stores) Replace the old password in the secret containing it with the new one. If the secret is a file in a file system volume secret store, ensure that the new password is encoded appropriately.

For example, base64-encode the password, and add it to the file:

```
$ echo -n bmV3a2V5cGFzc3dvcnQ= > keystoreA_keypass
```

4. List the keys and password strings contained in the keystore:

```
$ keytool -list -storetype JCEKS -keystore /path/to/openam/security/keystores/am_keystore.jceks>
```

5. Use the `keytool` command to change the password of each of the key aliases, for example:

AM Keystore

```
$ keytool -keypasswd -storetype JCEKS -keystore /path/to/openam/security/keystores/am_keystore.jceks -
alias mykey
Enter keystore password: Enter the password in the .am_keystore_storepass file
New key password for <mykey> Enter the password in the .am_keystore_keypass file
Re-enter new key password for <mykey> Enter the password in the .am_keystore_keypass file

** Caution **

Remember to change the passwords of the configstorepwd and the dsameuserpwd aliases.
Failure to do so will render AM unbootable.
```

Secret Stores

```
$ keytool -keypasswd -storetype JCEKS -keystore /path/to/openam/security/keystores/keystoreA.jceks -
alias mykey
Enter keystore password: Enter the password in the keystoreA_storepass file
New key password for <mykey> Enter the password in the keystoreA_keypass file
Re-enter new key password for <mykey> Enter the password in the keystoreA_keypass file

** Remember **

Secrets in file system volume secret stores are, by default, not in cleartext.
You need to decode them before using them with the keytool command.
```

6. (Optional) If you also need to change the keystore password, see "To Change the Keystore Password".

7. Ensure that password files and keystores are maintained on every instance in your environment. Every AM instance has its own keystores and password files.
8. (AM keystore) Restart the AM instances affected by the configuration changes.

Changing Keystore Passwords

Decrypting and viewing the contents of a keystore requires a password. This password is specified by the user at the time the keystore is created, but you might need to update the password at a later time.

To Change the Keystore Password

1. (AM keystore) Replace the old password in the `.am_keystore_storepass` file with the new one:

```
$ echo -n newpassword > /path/to/openam/security/secrets/default/.am_keystore_storepass
```

Important

Use `echo -n` to avoid inserting hidden trailing newline characters. Even if the `keytool` command is able to use the password in the file, AM may not be able to use the key aliases if there are hidden trailing newline characters in the password file.

2. (Secret stores) Replace the old password in the secret containing it with the new one. If the secret is a file in a file system volume secret store, ensure that the new password is encoded appropriately.

For example, base64-encode the password, and add it to the file:

```
$ echo -n bmV3c3RvcnVwYXNzd29yZA== > keystoreA_storepass
```

3. Change the password of the keystore:

AM Keystore

```
$ keytool -storepasswd -storetype JCEKS -keystore /path/to/openam/security/kestores/
am_keystore.jceks
Enter keystore password: Enter the password in the .am_keystore_storepass file.
New keystore password: Enter the new password.
Re-enter new keystore password:
```

Secret Stores

```
$ keytool -storepasswd -storetype JCEKS -keystore /path/to/openam/security/keystores/keystoreA.jceks
Enter keystore password: Enter the password in the keystoreA_storepass file.
New keystore password: Enter the new password.
Re-enter new keystore password:
```

```
** Remember **
```

Secrets in file system volume secret stores are, by default, *not* in cleartext.
That means, you need to decode them before using them with the keytool command.

4. (Optional) If you also need to change the key aliases' password, see "To Change Key Alias Passwords".
5. Ensure that password files and keystores are maintained on every instance in your environment. Each AM instance has its own keystores and password files.
6. (AM keystore only) Restart the AM instance or instances affected by the configuration changes.

Configuring Secret Stores

Secret stores are repositories for cryptographic keys and credentials. You can configure them globally, which ensures that all realms inherit your secret store settings. You can also configure secret stores by realm, which allows you to set different secret store settings for each realm.

Since secrets must be shared by all the servers in the site, a good practice is to keep them all under the same directory or mount point, for example, `/path/to/openam/security/secrets`.

AM provides a keystore and a filesystem secret store by default in new installations, but we recommend that you create your own secret stores in production environments.

AM supports the following secret store types:

- Environment and System Properties

AM supports configuring secrets derived from JVM system properties.

- Keystore

AM supports a number of different keystore formats, including JCEKS, JKS, PKCS11, and PKCS12.

- File System Secret Volumes

AM supports secrets that are stored as files in defined folders. For example, in a cloud deployment you could mount a secret volume that AM can access.

- Hardware Security Modules (HSM)

AM supports retrieval of secrets from hardware security modules, either locally or over the network.

- Google Cloud Key Management Service (KMS)

AM supports retrieving secrets from the Google Cloud Platform KMS.

Tasks to Configure Secret Stores

Task	Resources
<p>Understand How AM Resolves Secrets</p> <p>Secrets are first resolved at the realm level, and then globally.</p>	<ul style="list-style-type: none"> • "Understanding How AM Resolves Secrets"
<p>Configure Secret Stores</p> <p>Configure as many secret stores as your environment needs.</p>	<ul style="list-style-type: none"> • "The Environment and System Property Secrets Store" • "Keystore Secret Stores" • "File System Secret Volumes Secret Stores" • "HSM Secret Stores" • "Google KMS Secret Stores"
<p>Map Secret IDs to Secrets</p> <p>A number of AM features require the use of secrets for signing and encryption. For each requirement, AM has a secret ID.</p> <p>You can create active aliases in keystore and HSM secret stores.</p>	<ul style="list-style-type: none"> • "Mapping and Rotating Secrets"

Understanding How AM Resolves Secrets

Most secret stores are configured at the global level, by going to **Configure > Secret Stores**, or at the realm level, by going to **Realms > *Realm Name* > Secret Stores**.

Secrets derived from environmental or system properties are configured globally, in a special, persistent secret store.

When resolving secrets, AM will search secret stores in the following order:

1. Any secret store configured for the realm, regardless of their type.
2. Any secret store configured globally, regardless of their type.

If AM cannot find the alias, it will log an error and the operation it was trying to do (for example, signing a client-based session token) will fail.

Caution

Map each secret ID *once* across the secret stores configured for the realm, or globally. For example, in a realm with two secret stores configured (a keystore secret store and a HSM secret store) the `am.services.oauth2.jwt.authenticity.signing` secret ID is mapped only in the keystore secret store and not in the HSM secret store.

The Environment and System Property Secrets Store

There is a global instance of the Environment and System Property Secrets Store configured at all times. You can access the Environment and System Property Secrets Store globally.

Secrets within the environment and system property secrets store are derived from system properties with the same key as the secret value name (for example, `am.services.oauth2.stateless.token.encryption`), or as environment variables with keys that have the secret value name in upper case, and with all period characters replaced with underscores (for example, `AM_SERVICES_OAUTH2_STATELESS_TOKEN_ENCRYPTION`).

AM configures this secret store on startup. Restart AM or the container where it runs to use the new secret mappings.

The only configuration settings that apply for the environment and system property secrets store is the format of the secrets. Secrets that come from this store cannot be rotated, retired (deleted), or removed.

To Configure the Environment and System Property Secrets Store

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Go to Configure > Secret Stores > Environment and System Property Secrets Store.
3. From the Value format drop-down list, select one of the following:
 - Plain Text: the secret is provided in UTF-8 encoded text.
 - Base64 encoded: the secret is provided in Base64 encoded binary values.
 - Encrypted text: the plain text secrets are encrypted using AM's encryption key, found at Deployment > Servers > Security > Encryption.
 - Encrypted Base64 encoded: the Base64 encoded binary values are encrypted using AM's encryption key.
 - Encrypted HMAC key: the Base64 encoded binary representation of the HMAC key is encrypted using AM's encryption key.
 - BASE64_HMAC_KEY: the Base64 encoded binary representation of the HMAC key.
4. Save your changes.

Keystore Secret Stores

A keystore secret store is a secret store that maps to a keystore file, for example, a JKS, JCEKS, PKCS11, or PKCS12 file.

Tip

During installation or after an upgrade from a version of AM earlier than 6.5, AM deploys a number of secret stores. You can use them as an example to configure your own secret stores. For more information, see [About the Default Secret Stores](#).

To Create a Keystore Secret Store

Keystore secret stores can be configured at a global or realm level:

1. To create on a global level:

- Go to **Configure > Secret Stores**.

To create on a realm level:

- Go to **Realms > *Realm Name* > Secret Stores**.

2. Select **Add Secret Store**.
3. Enter the **Secret Store ID**.
4. From the **Store Type** drop-down list, select **Keystore**.
5. Enter the keystore file to use.

This file must be available to all AM instances, for example, by storing it on a shared filesystem, or by copying and maintaining the file across instances.

6. Select **Create**.

To Configure a Keystore Secret Store

1. To configure a global keystore:

- Go to **Configure > Secret Stores**.

To configure a realm keystore:

- Go to **Realms > *Realm Name* > Secret Stores**.

2. Select the store you want to modify.
3. Enter the keystore file name in the **File** field.
4. Enter the **Keystore Type**, for example **JKS**, **JCEKS**, **PKCS11**, or **PKCS12**.

The specified keystore type must be supported by, and configured in, the local Java runtime environment.

5. Set the Provider name. If blank, the JRE default will be used.
6. In the Store password secret ID field, enter the secret ID from which AM will resolve the password that opens the keystore file, or none if the password is blank. For example, `storepass`.

AM resolves this secret ID using the other secret stores configured. For example, a file system secret volume secret store mapped to the directory where the file containing the password is stored, or an HSM secret store. For more information about how AM resolves secrets, see ["Understanding How AM Resolves Secrets"](#).

7. In the Entry password secret ID field, enter the secret ID from which AM will resolve the password to the keys stored in the keystore, or none if the password is blank. For example, `entrypass`.

AM resolves this secret ID using the other secret stores configured. For example, a file system secret volume secret store mapped to the directory where the file containing the password is stored, or an HSM secret store. For more information about how AM resolves secrets, see ["Understanding How AM Resolves Secrets"](#).

8. Set the Key lease expiry time in minutes.
9. Save your changes.

File System Secret Volumes Secret Stores

File System Secret Volumes maps to a directory storing files that contain secrets - one secret per file. For a given secret value, file system secret volumes stores will look for a file with the same name as the secret value name, and read its contents using the configured value format. They can be configured at a global and realm level.

Tip

During installation or after an upgrade from a version of AM earlier than 6.5, AM deploys a number of secret stores. You can use them as an example to configure your own secret stores. For more information, see [About the Default Secret Stores](#).

To Create a File System Secret Volume Store

1. To create on a global level:
 - Go to Configure > Secret Stores.

To create on a realm level:

- Go to Realms > *Realm Name* > Secret Stores.

2. Select Add Secret Store.
3. Enter the Secret Store ID.
4. From the Store Type drop-down list, select File System Secret Volumes.
5. Enter the name of the directory containing the secret files.

This directory must be available to all AM instances, for example, by converting it to a shared filesystem, or by creating and maintaining it and its files across instances.

6. Select Create.

To Configure a File System Secret Volume Store

1. To configure a global file system secret volume store:

- Go to Configure > Secret Stores.

To configure a realm file system secret volume store:

- Go to Realms > *Realm Name* > Secret Stores.

2. Select the store you want to modify.
3. Enter the directory name in the Directory field.
4. (Optional) Enter a suffix to add to the name of each secret in the File suffix field. For example, `txt`.
5. From the Value format drop-down list, select one of the following:
 - Plain Text: the secret is provided in UTF-8 encoded text.
 - Base64 encoded: the secret is provided in base64-encoded binary values.
 - Encrypted text: the plain text secrets are encrypted using AM's encryption key, found at Deployment > Servers > Security > Encryption.
 - Encrypted Base64 encoded: the base64-encoded binary values are encrypted using AM's encryption key.
 - Encrypted HMAC key: the base64-encoded binary representation of the HMAC key is encrypted using AM's encryption key.
 - Base64 encoded HMAC key: the base64-encoded binary representation of the HMAC key.
 - Encrypted with Google KMS: the secrets are encrypted with a secret stored in the Google Cloud KMS, then base64-encoded.

See "Using Google Cloud KMS Secrets to Decrypt AM Secrets".

- Google KMS-encrypted HMAC key: the HMAC key is encrypted with a secret stored in the Google Cloud KMS, then base64-encoded.

See "Using Google Cloud KMS Secrets to Decrypt AM Secrets".

6. Save your changes.

You can now map secret IDs to files stored in the secret store directory. See "To Map Files in File System Secret Volumes Secret Stores".

HSM Secret Stores

An HSM Secret Store maps to a hardware security module. To configure an HSM Secret Store, you need a secret ID that can provide the PIN or password for the HSM, or an extension can be created that provides a Guice binding for a custom PKCS11 `java.security.Provider` to obtain the keystore.

To Create an HSM Secret Store

HSM Secret Stores can be configured at a global and realm level:

1. To create on a global level:

- Go to Configure > Secret Stores.

To create on a realm level:

- Go to Realms > *Realm Name* > Secret Stores.

2. Select Add Secret Store.

3. Enter the Secret Store ID.

4. From the Store Type drop-down list, select HSM.

5. Enter the Configuration File containing initialization configuration for the HSM.

6. In the Provider Guice Key Name field, enter the name of a Guice key that can be used to obtain an initialized provider from which the HSM keystore can be obtained.

7. In the HSM PIN/password secret ID field, enter the secret ID from which HSM's PIN or password can be obtained.

AM resolves this secret ID using the other secret stores configured. For example, a file system secret volume secret store mapped to the directory where the file containing the password is stored, or a keystore secret store. For more information about how AM resolves secrets, see "Understanding How AM Resolves Secrets".

8. Select Create.

To Configure an HSM Secret Store

1. To configure a global HSM secret store store:
 - Go to Configure > Secret Stores.

To configure a realm HSM secret store:

 - Go to Realms > *Realm Name* > Secret Stores.
2. Select the store you want to modify.
3. In the Configuration File field, enter the name of the file containing initialization configuration for the HSM.
4. In the Provider Guice Key Name field, enter the name of a Guice key that can be used to obtain an initialized provider from which the HSM keystore can be obtained.
5. In the HSM PIN/password secret ID field, enter the secret ID from which HSM's PIN or password can be obtained.

AM resolves this secret ID using the other secret stores configured. For example, a file system secret volume secret store mapped to the directory where the file containing the password is stored, or a keystore secret store. For more information about how AM resolves secrets, see "Understanding How AM Resolves Secrets".
6. Set the Key lease expiry time in minutes.
7. Save your changes.

Google KMS Secret Stores

You can configure AM to retrieve secrets from the Google Cloud KMS. Support includes:

- Mapping Google Cloud KMS secrets to secret IDs used for signing and verification purposes. Using Google Cloud KMS secrets as mappings for encryption and decryption secret IDs is *not* supported.

For example, mapping a Google Cloud KMS secret to the `am.services.oauth2.oidc.signing.RSA` secret ID is supported because it is a secret ID used for signing OAuth 2.0 tokens. Mapping a Google Cloud KMS secret to the `am.services.oauth2.oidc.decryption.RSA1.5` secret ID is *not* supported because it is used for decrypting OpenID Connect parameters.

+ *Supported Signing Algorithms for Google Cloud KMS Secrets*

```
SHA256WithRSA (RS256)
SHA512WithRSA (RS512)
SHA256WithRSAAndMGF1 (PS256)
SHA512WithRSAAndMGF1 (PS512)
```

```
SHA256WithECDSA (ES256)
SHA384WithECDSA (ES384)
```

Caution

Signing tokens with Google Cloud KMS secrets is not a fast operation. For every signature request, AM makes an API call to the Google Cloud KMS to perform the signature operation.

Test the time it would take in your environment to sign tokens under stress conditions to determine if the delay is acceptable. We recommend that you use Google Cloud KMS secrets in environments with a low volume of signatures and high volume of verifications, since AM performs the verification locally.

- Using a Google Cloud KMS secret to decrypt secrets loaded using other secret stores, or to decrypt the hashed password of the `amAdmin` user.

Prerequisites

You need a Google Cloud Platform account that has a project. The project must have:

- A key ring containing the secrets that AM will use. It can be configured in any Google Cloud location.
- A service account that AM will use to connect to the project.

Refer to the [Google Key Management Service documentation](#) and Google's [Getting Started with Authentication](#) for more information.

Configuring the Google Service Account Credentials

On a Google Cloud environment, AM uses Google's Java SDK to communicate with the Google Cloud KMS directly. This means that, as long as your Google Cloud environment has a default service account, AM will use it automatically.

If you do not have a default service account or do not want to use it for this purpose, or if you are using Google Cloud KMS secret stores in a non-Google Cloud environment, you must configure the path to the credentials in an environment variable so that AM can use them:

1. Log in to your Google Cloud Platform Account.
2. Download the credentials file for the Google service account that AM will use to connect to the project, and store it in the server where AM runs.
3. Set up the `GOOGLE_APPLICATION_CREDENTIALS` environment variable to the path of the credentials. Ensure that the variable is available to the container where AM runs.

For example, add the environment variable to the `setenv.sh` file of your Apache Tomcat installation:

```
export GOOGLE_APPLICATION_CREDENTIALS="/path/to/Tomcat/Google-service-account-credentials-for-AM.json"
```

4. Restart the container where AM runs.
5. Perform the steps in this procedure on each of the servers where AM runs.

Creating KMS Secret Stores

Google KMS secret stores can be configured at a global or realm level:

1. To create on a global level:

- Go to Configure > Secret Stores.

To create on a realm level:

- Go to Realms > *Realm Name* > Secret Stores.

2. Select Add Secret Store.
3. Enter the Secret Store ID.
4. From the Store Type drop-down list, select Google KMS.
5. In the Project field, enter the Google Cloud Platform project that contains the key ring with the secrets.

At the time of this writing, you can find your projects by logging in to your Google Cloud Platform dashboard.

6. Configure the following fields related to the key ring.

At the time of this writing, you can find the required information by logging in to the Google Cloud Platform dashboard, choosing your project, and then going to > Security > Cryptographic Keys.

- a. In the Location field, enter the location of the key ring.
- b. In the Key Ring field, enter the name of the key ring containing the secrets that AM should use.

7. Select Create.

The page of the new secret store appears.

8. (Optional) Configure the size of the public key cache and its duration as required in your environment.

+ *Tips and Notes About the Public Key Cache*

When AM signs data with a secret stored in the Google Cloud KMS, it makes an API call to the Google Cloud KMS to perform the signature operation.

When AM needs to verify a signature, it retrieves the public key from the Google Cloud KMS and verifies the signature locally. The cache prevents AM from retrieving the public key every time, and therefore, speeds the verification process.

The cache lives in AM's heap and it is created on each of the AM instances for each of the Google Cloud KMS secret stores. We recommend that you leave the default settings unless you have a large number of keys in a key chain.

Setting a long cache timeout may be more efficient, since AM does not need to contact the Google Cloud KMS to retrieve public keys that often, but note that AM will not detect if you have marked a key as expired in the Google Cloud KMS until the cache expires.

Using Google Cloud KMS Secrets to Decrypt AM Secrets

You can use a Google Cloud KMS secret to decrypt secrets stored in AM secret stores as they are read from the filesystem, environment variables, or system properties.

You can also use the same secret to decrypt the hashed password of the `amAdmin` user. See "Changing the amAdmin Password (Secret Stores)".

Important

You can only configure one Google Cloud KMS secret for decrypting secrets in the AM site.

This procedure assumes that the encrypted secrets will be stored in a filesystem, and therefore, configured in AM in a file system volume secret store:

1. Check if you already have a Google Cloud KMS secret for decrypting.

Go to **Configure > Server Defaults > Advanced**, and check if the `org.forgerock.openam.secrets.googlekms.decryptionkey` advanced server property is configured.

If it is, you do not need to create another key.

If the property is not configured, log in to your Google Cloud dashboard and create a secret of one of the following types in the key ring of your choosing:

- Symmetric encrypt/decrypt
- Asymmetric decrypt

2. Use the secret you identified or created in the previous step to encrypt the secrets that AM will use.

You can use the **gcloud** tool included in Google Cloud's SDK to encrypt the secrets. The tool creates a binary file with the encrypted secret, but AM does not support secrets in binary format. To work around this, base64-encode the encrypted secret. For example:

```
gcloud kms encrypt \
--plaintext-file=./secret.txt \
--ciphertext-file=- \
--project=my_project_ID \
--location=my_location \
--keyring=my_keyring_for_AM \
--key=my_key_for_decrypting_secrets_in_AM | base64 > secret.enc
```

3. Rename the files containing the secrets so that they map to the required secret IDs. Use the tables in "Secret ID Default Mappings" for guidance.

For example, to create a mapping for the Web and Java agents' OAuth 2.0 provider, rename the file containing the relevant secret to a file called `am.global.services.oauth2.oidc.agent.idtoken.signing`.

Depending on the configuration of the secret store, you may be able to add a suffix to the file name, such as `.enc`.

4. Share the encrypted secrets with the AM servers. This may mean, for example, copying the encrypted files to the same directory in every AM server, or mounting a directory in every AM server that is shared across the instances.
5. In the AM console, go to Configure > Server Defaults > Advanced.
6. (Optional) If unset, set the `org.forgerock.openam.secrets.googlekms.decryptionkey` advanced server property to the fully qualified resource ID of the Google Cloud KMS secret that you used in the previous step. For example:

```
projects/my_project_ID/locations/my_location/keyRings/my_keyring_for_AM/cryptoKeys/
my_key_for_decrypting_secrets_in_AM
```

For information about how to find the key ID, see [Object Hierarchy](#) in the Google Cloud KMS documentation.

7. Configure the file system volume secret store that points to the directory containing the encrypted secrets. See "To Configure a File System Secret Volume Store".

Mapping and Rotating Secrets

Several AM features require the use of secrets for signing and encryption. For each requirement, AM has a **secret ID**. To provide AM with the required secret, map one or more aliases from the secret stores you configure to each of the secret IDs. These mappings allow you to choose which is the active aliases, and rotate them when they become expired or compromised.

For a list of secret IDs and their default mappings, see "Secret ID Default Mappings".

Active aliases are used for signature generation and encryption, while the non-active aliases are used for signature verification and decryption. An non-active alias can be rotated to become the active alias, while the old alias will still remain valid. Non-active secrets are mainly used for signature verification and decryption. A secret can be retired when it is no longer considered secure.

For example, if you mapped several aliases for signing OAuth 2.0 client-based tokens, new tokens are signed with the active secret, and incoming tokens are verified against both the active and the non-active secrets.

To Map Aliases in Keystore, HSM, or Google KMS Secret Stores

1. To map secrets within a global secret store:

- Navigate to Configure > Secret Stores.

To map secrets within a realm secret store:

- Navigate to Realms > *Realm Name* > Secret Stores.

2. Select the store that contains the secrets you want to map.
3. On the Mappings tab, select Add Mapping.
4. From the Secret ID drop-down list, select the Secret ID that is to be associated to an alias.

For information about the different secret ID mappings, see "Secret ID Default Mappings".

5. Enter any Alias and click the add (+) icon.

You can add as many aliases as necessary. The first alias in the list determines which alias is the active one. Active secrets are used for signature generation and encryption, while the non-active secrets are used for signature verification and decryption.

Tip

When configuring mappings for a Google KMS secret store, map one secret for each secret ID, and manage key rotation in the Google Cloud KMS key ring.

6. Drag and drop to change the order of aliases, and set which alias is active.
7. If an alias is considered no longer secure, it can be retired by clicking the delete (✕) icon.
8. Once your mappings are complete, select Create.

To Map Files in File System Secret Volumes Secret Stores

File system secret volumes secret stores do not allow rotating or retiring secrets through mappings like other stores do.

To map secret IDs to files, perform the following steps:

1. Change paths to the directory configured in the secret store. For example, change paths to `/openam/secrets`.

2. Create the required files to store your secrets using the tables in "Secret ID Default Mappings" for guidance, and leave them empty.

For example, to create a mapping for the Web and Java agents' OAuth 2.0 provider, create a file called `am.global.services.oauth2.oidc.agent.idtoken.signing`.

You may also create mappings for secret store-specific secrets, such as the keystore secret store password, the keystore secret store entry password, or the HSM guice key. These mappings do not require specific secret IDs. For example, you can create a file called `mykeystorepassword`, and then configure it in the Store password secret ID field of your keystore secret store.

The name of a secret ID—and therefore the file names given to file system secrets—must consist of only alphanumeric characters and periods(.). The names cannot start or end with periods, or have more than one period in a row.

Depending on the configuration of the secret store, you may be able to add a suffix to the file name, such as `.txt`.

3. Store the relevant secret value in each file.

The format of the secret value depends on the configuration of the secret store. For example, if you have configured File Format to be `Encrypted text`, you must encode the secret value with AM's encryption key.

+ *How Do I Encode Secrets With AM's Encryption Key?*

Use the `https://openam.example.com:8443/openam/encode.jsp` page to encode the secret, then add the encoded value to the secret file.

Ensure that secrets do not contain trailing new line characters. If you are using the `echo` command to add secrets to a file, append the `-n` option. For example:

```
$ echo -n AQICmX1ntZv3XETMgDo+0zFynC8UMGJgop+K \  
> am.global.services.oauth2.oidc.agent.idtoken.signing
```

Secret ID Default Mappings

The following groups contain the secret IDs used by the AM features, and their default mappings, if any. Expand the categories for additional information about where or how the mappings are used.

OAuth 2.0 and OpenID Connect

+ *Secret ID Mappings for JWT Authenticity Signing*

The following table shows the secret ID mapping used to sign several OAuth 2.0 and OpenID Connect-related JWTs:

Secret ID	Default Alias	Algorithms
am.services.oauth2.jwt.authenticity.signing	hmacsigningtest	HS256 HS384 HS512

Note

This key is used to sign the following tokens and requests:

- OpenID Connect tokens for Web and Java Agents.
- OpenID Connect tokens that are signed with an HMAC algorithm.
- Macaroon access and refresh tokens.
- Consent requests to remote consent agents that are signed with an HMAC algorithm.

+ Secret ID Mappings for Encrypting Client-Based OAuth 2.0 Tokens

The following table shows the secret ID mapping used to encrypt client-based access tokens:

Secret ID	Default Alias	Algorithms
am.services.oauth2.stateless.token.encryption	directentest	A128CBC-HS256

+ Secret ID Mappings for Signing Client-Based OAuth 2.0 Tokens

The following table shows the secret ID mappings used to sign client-based access tokens:

Secret ID	Default Alias	Algorithms
am.services.oauth2.stateless.signing.ES256	es256test	ES256
am.services.oauth2.stateless.signing.ES384	es384test	ES384
am.services.oauth2.stateless.signing.ES512	es512test	ES512
am.services.oauth2.stateless.signing.HMAC	hmacsigningtest	HS256 HS384 HS512
am.services.oauth2.stateless.signing.RSA	rsajwtsginingkey	PS256 PS384 PS512 RS256 RS384

Secret ID	Default Alias	Algorithms
		RS512

+ Secret ID Mappings for Signing Remote Consent Requests

The following table shows the secret ID mappings used to sign remote consent requests:

Secret ID	Default Alias	Algorithms ^a
am.applications.agents.remote.consent.request.signing.ES256	es256test	ES256
am.applications.agents.remote.consent.request.signing.ES384	es384test	ES384
am.applications.agents.remote.consent.request.signing.ES512	es512test	ES512
am.applications.agents.remote.consent.request.signing.RSA	rsajwtsigningkey	RS256 RS384 RS512 PS256 PS384 PS512

^a If you select an HMAC algorithm for signing consent requests (**HS256**, **HS384**, or **HS512**), the value of the Remote Consent Service secret property is used, instead of an entry from the secret stores.

Since the HMAC secret is shared between AM and the remote consent client, a malicious user compromising the client could potentially create tokens that AM would trust. Therefore, to protect against misuse, AM also signs the token using a non-shared signing key configured in the **am.services.oauth2.jwt.authenticity.signing** secret ID.

+ Secret ID Mappings for Decrypting Remote Consent Responses

The following table shows the secret ID mapping used to decrypt remote consent responses:

Secret ID	Default Alias	Algorithms ^a
am.services.oauth2.remote.consent.response.decryptiontest		RSA-OAEP-256

^a If you select an algorithm other than **RSA-OAEP-256** for decrypting consent responses, the value of the Remote Consent Service secret property is used, instead of an entry from the secret stores.

+ Secret ID Mappings for Decrypting OpenID Connect Request Parameters

The following table shows the secret ID mapping used to decrypt OpenID Connect request parameters:

Secret ID	Default Alias	Algorithms ^a
am.services.oauth2.oidc.decryption.RSA1.5	test	RSA with PKCS#1 v1.5 padding

Secret ID	Default Alias	Algorithms ^a
am.services.oauth2.oidc.decryption.RSA.OAEP	test	RSA with OAEP with SHA-1 and MGF-1
am.services.oauth2.oidc.decryption.RSA.OAEP.256	test	RSA with OAEP with SHA-256 and MGF-1

^a The following applies to confidential clients only:

If you select an AES algorithm (**A128KW**, **A192KW**, or **A256KW**) or the direct encryption algorithm (**dir**), the value of the Client Secret field in the OAuth 2.0 Client is used as the secret instead of an entry from the secret stores.

The following signing and encryption algorithms use the Client Secret field to store the secret:

- Signing ID tokens with an HMAC algorithm
- Encrypting ID tokens with AES or direct encryption
- Encrypting parameters with AES or direct encryption

Store only *one* secret in the Client Secret field; AM will use different mechanisms to sign and encrypt depending on the algorithm, as explained in the OpenID Connect Core 1.0 errata set 1 specification.

+ Secret ID Mappings for Signing OpenID Connect Tokens

The following table shows the secret ID mapping used to sign OpenID Connect ID tokens:

Secret ID	Default Alias	Algorithms ^a
am.services.oauth2.oidc.signing.ES256	es256test	ES256
am.services.oauth2.oidc.signing.ES384	es384test	ES384
am.services.oauth2.oidc.signing.ES512	es512test	ES512
am.services.oauth2.oidc.signing.RSA	rsajwt signingkey	PS256 PS384 PS512 RS256 RS384 RS512
am.services.oauth2.oidc.signing.EDDSA	—	EdDSA with SHA-512

^a The following applies to confidential clients only:

If you select an HMAC algorithm for signing ID tokens (**HS256**, **HS384**, or **HS512**), the Client Secret property value in the OAuth 2.0 Client is used as the HMAC secret instead of an entry from the secret stores.

Since the HMAC secret is shared between AM and the client, a malicious user compromising the client could potentially create tokens that AM would trust. Therefore, to protect against misuse, AM also signs the token using a non-shared signing key configured in the **am.services.oauth2.jwt.authenticity.signing** secret ID.

+ Secret ID Mappings for CA Certificates Used in mTLS Client Authentication

The following table shows the secret ID mapping used to store the CA certificates AM should trust during mTLS client authentication:

Secret ID	Default Alias	Algorithms
am.services.oauth2.tls.client.cert.authentication	—	—

Web Agents and Java Agents

+ Secret ID Mappings for the Agents' OAuth 2.0 Provider

The following table shows the secret ID mapping used sign the JWTs provided to Web and Java agents:

Secret ID	Default Alias	Algorithms
am.global.services.oauth2.oidc.agent.idtoken.signing	rsajwt signingkey	RS256 RS384 RS512

Authentication

+ Secret ID Mappings for Persistent Cookies

The following table shows the secret ID mappings used to encrypt and then sign persistent cookies:

Secret ID	Default Alias	Algorithms
am.default.authentication.modules.persistentcookie.encryption	test	RSA (at least 2048 bits)
am.default.authentication.modules.persistentcookie.signing	hmacsigningtest	HS256

For each instance of a persistent cookie module available in a realm, there is a dynamic secret ID associated with that module configuration instance.

For example, in a single realm you can have a Persistent Cookie module instance with the name *helloworld*, and a separate Persistent Cookie module instance with the name *hellomars*.

The following secret ID mappings could be used to encrypt and then sign persistent cookies:

Secret ID	Default Alias	
am.authentication.modules.persistentcookie. <i>helloworld</i> .encryption	helloworld	
am.authentication.modules.persistentcookie. <i>helloworld</i> .signing	hmacsigninghelloworld	

Secret ID	Default Alias	
am.authentication.modules.persistentcookie.hellomars.encryption	hellomars	
am.authentication.modules.persistentcookie.hellomars.signing	hmacsigninghellomars	

AM will attempt to look up the secrets with the Persistent Cookie module instance name. If unsuccessful, AM will look up the secrets using the default secret ID.

+ Secret ID Mappings for Encrypting Authentication Trees' Secure State Data

The following table shows the secret ID mapping used to encrypt sensitive data stored in the secure state of an authentication tree:

Secret ID	Default Alias	Algorithms
am.authn.trees.transientstate.encryption	directentest	AES 256-bit

SAML v2.0

+ Secret ID Mappings for Encrypting SAML v2.0 Local Storage JWTs

The following table shows the secret ID mapping used to encrypt the JWTs SAML v2.0 creates in local storage:

Secret ID	Default Alias	Algorithms
am.global.services.saml2.client.storage.jwt.encryption	directentest	A256GCM

+ Secret ID Mappings for Signing SAML v2.0 Metadata

The following table shows the secret ID mappings used to sign SAML v2.0 metadata:

Secret ID	Default Alias	Algorithms
am.services.saml2.metadata.signing.RSA	rsajwtssigningkey	RSA SHA-256

+ Secret ID Mappings for SAML v2.0 Signing and Encryption

The following table shows the secret ID mappings used to sign and encrypt SAML v2.0 elements:

Secret ID	Default Alias	Algorithms
am.default.applications.federation.entity.providers.saml2.idp.encryption	directentest	RSA with PKCS#1 v1.5 padding

Secret ID	Default Alias	Algorithms
am.default.applications.federation.entity.providers.saml2.idp.signing	rsa-jwt-signing-key	RSA with OAEP
am.default.applications.federation.entity.providers.saml2.idp.signing	rsa-jwt-signing-key	RSA SHA-1 ^a ECDSA SHA-256 ECDSA SHA-384 ECDSA SHA-512 RSA SHA-256 RSA SHA-384 RSA SHA-512 DSA SHA-256
am.default.applications.federation.entity.providers.saml2.sp.encryption	rsa-jwt-signing-key	RSA with PKCS#1 v1.5 padding RSA with OAEP
am.default.applications.federation.entity.providers.saml2.sp.signing	rsa-jwt-signing-key	RSA SHA-1 ^a ECDSA SHA-256 ECDSA SHA-384 ECDSA SHA-512 RSA SHA-256 RSA SHA-384 RSA SHA-512 DSA SHA-256

^aThis algorithm is for compatibility purposes only, and its use should be avoided.

You can specify a custom secret ID identifier for each hosted SAML v2.0 entity provider in a realm, which creates new secret IDs. These secret IDs can be unique to the provider, or shared by multiple providers.

For example, you could add a custom secret ID identifier named *mySamlSecrets* to a hosted identity provider.

AM dynamically creates the following secret IDs, which the hosted identity provider uses for signing and encryption:

- `am.applications.federation.entity.providers.saml2.mySamlSecrets.signing`
- `am.applications.federation.entity.providers.saml2.mySamlSecrets.encryption`

AM will attempt to look up the secrets with the custom secret ID identifier. If unsuccessful, AM will look up the secrets using the default secret IDs.

IoT

+ Secret ID Mappings for the IoT Trusted JWT Issuer

The following table shows the secret ID mapping that the IoT service uses when configured as a trusted OAuth 2.0 JWT issuer:

Secret ID	Default Alias	Algorithms
am.services.iot.jwt.issuer.signing	hmacsigningtest	HS256

+ Secret ID Mappings for IoT Certificate Verification

The following table shows the secret ID mapping for the CA certificate that the "Register Thing Node" in the *Authentication and Single Sign-On Guide* uses to verify the X.509 digital certificate included in the proof-of-possession JWT.

Secret ID	Default Alias	Algorithms
am.services.iot.cert.verification	—	—

Changing Default Key Aliases

For demo and test purposes, AM configures different [demo key aliases](#) for several features. You can keep the demo key aliases configured in those features you are not using, but you may decide to remove them completely from your production environment.

When possible, the following list includes the Global Services or Server Default paths where the demo key aliases are configured. If you already have configured any of the features in a realm, ensure that the key alias is replaced in the realm configuration as well.

To replace the default key aliases:

1. Create the required key aliases following the tasks in "Managing Key Aliases and Passwords".
2. Change default key aliases in AM:

Web Agents and Java Agents

Refer to the *ForgeRock Web Agents User Guide* and the *ForgeRock Java Agents User Guide* for more information.

Persistent Cookie Module

To change the default mapping for the Persistent Cookie module, go to Realms > *Realm Name* > Authentication > Settings > Security and replace the **test** key alias in the Persistent Cookie Encryption Certificate Alias field with the alias you created for persistent cookies in your secret stores.

For more information about the secret ID mappings used by this feature, see [Secret ID Mappings for Persistent Cookies](#).

OAuth 2.0 and OpenID Connect Providers

See the list of secret IDs and their defaults [here](#).

SAML v2.0 Hosted Providers

See the list of secret IDs and their defaults [here](#).

Client-Based Sessions

Go to Configure > Global Services > Session > Client-based Sessions and replace the **test** key alias in the Signing RSA/ECDSA Certificate Alias field and in the Encryption RSA Certificate Alias field.

User Self-Service

Go to Realms > *Realm Name* > Services > User Self-Service and populate the values of the Encryption Key Pair Alias and the Signing Secret Key Alias properties.

Note that the name of the demo keys shows with a gray color; that does not mean the fields are filled in.

Authentication Trees

Authentication trees use the secret ID specified in Secret ID Mappings for Encrypting Authentication Trees' Secure State Data.

Important

Ensure that this secret ID is always mapped to an existing, resolvable secret or key alias, or authentication trees may not work as expected.

IoT

The IoT Service uses the secret IDs specified in Secret ID Mappings for the IoT Trusted JWT Issuer.

Chapter 6

Securing the Session Cookie

After authenticating an end user, AM stores their session (for client-based sessions), or a pointer to their session (for CTS-based sessions), in a cookie in the end user's browser.

HTTPS communication already helps to keep cookies secure since the encrypted communication cannot be eavesdropped. However, there are other ways a malicious user can hijack a cookie. For example, cross-site scripting (XSS) and cross-site tracing (XST) involve injecting HTML or JavaScript on a legitimate website. By using JavaScript code, the malicious user can steal the cookie directly from the browser.

The following table summarizes the tasks you need to perform to protect session cookies:

Task	Resources
<p>Configure the HttpOnly Flag</p> <p>This flag ensures that the session cookie is transmitted over an HTTP or HTTPS channel only, protecting your environment against most XSS attacks.</p>	"Configuring HttpOnly Session Cookies"
<p>Configure the secure Flag</p> <p>This flag ensures the session cookie is only transmitted over HTTPS channels such that the session cookie is not carried over insecure HTTP redirections.</p>	"Configuring Secure Session Cookies"
<p>Choose a Session Cookie Name</p> <p>Change the name of the session cookie from the default of iPlanetDirectoryPro.</p>	"Changing the name of the Session Cookie"
<p>Restrict CDSSO Tokens to protect them against hijacking</p> <p>By default, AM provides a CDSSO tokens valid for the appropriated realms. Restrict tokens so that AM issues different tokens for different realms.</p>	"Enabling Restricted Tokens for CDSSO Session Cookies"
<p>Use host-only cookies</p> <p>Because host-only cookies are more secure than domain cookies, you <i>should</i> use host-only cookies unless you have a good business case for using domain cookies.</p>	Cookie Domains in the <i>Reference</i>

Tip

Client-based sessions are more vulnerable to hijacking, since they contain all the session information. To configure additional security measures, see "Configuring Client-Based Session Security".

Configuring HttpOnly Session Cookies

Whether you use HTTP or HTTPS, flag your cookies as **HttpOnly**, which means they are transmitted only over HTTP or HTTPS protocols. This setting alone already prevents most XSS attacks, since **HttpOnly** cookies cannot be transmitted using JavaScript.

Important

When a client makes a call to the `/json/authenticate` endpoint appending a valid SSO token, AM returns the `tokenId` field **empty** when **HttpOnly** cookies are enabled. For example:

```
{
  "tokenId": "",
  "successUrl": "/openam/console",
  "realm": "/"
}
```

To Configure the HttpOnly Flag

1. Log in to the AM console as an administrative user, for example, `amAdmin`.
2. Navigate to Configure > Server Defaults > Advanced.
3. Set the `com.sun.identity.cookie.httponly` advanced server property to `true`, and save your changes.

You must make this change in all the AM instances on the site.

Note

Regardless of the value of the `com.sun.identity.cookie.httponly` property, AM upgrades cookies to secure cookies (except the `amlbcookie` cookie) when requests arrive over a secure channel.

4. Restart AM or the container where it runs.

Configuring Secure Session Cookies

When using HTTPS, mark all your cookies as secure, which means they are only transmitted over HTTPS protocols.

This flag is useful for sites that allow both HTTPS and HTTP traffic, since it protects from HTTP redirection carrying session cookies across unencrypted connections.

To Configure the Secure Flag

1. Log in to the AM console as an administrative user, for example, `amAdmin`.
2. Navigate to Configure > Server Defaults > Security > Cookie.
3. Enable the Secure Cookie switch, and save your changes.
4. Restart AM or the container where it runs.

Changing the name of the Session Cookie

By default, the session cookie name is `iPlanetDirectoryPro`.

You must change this value to something unique in your environment that does not give away its contents. Do not use names such as `sessionCookie`.

Caution

If you change the name of the cookie on a production system, you are invalidating the sessions of any user that still had a valid cookie.

To Change the Name of the Session Cookie

1. Log in to the AM console as an administrative user, for example, `amAdmin`.
2. Navigate to Configure > Server Defaults > Security > Cookie.
3. Change the name in the Cookie Name field and save your changes.
4. Restart AM or the container where it runs.

Note that Web agents need to know the name of the session cookie. You must change this configuration when you change the name of the session cookie in AM. For more information, check the `com.sun.identity.agents.config.cookie.name` bootstrap property in the *ForgeRock Web Agents User Guide*.

Enabling Restricted Tokens for CDSSO Session Cookies

When the session cookie is a cross-domain single-sign on (CDSSO) cookie, meaning that it is valid across several domains, the damage a malicious user can cause is increased.

A malicious user who steals a CDSSO cookie can potentially use it to access any realms that session has logged into, which may span multiple domains. For example, a token stolen from `myapp.example.com` could be used to access `payroll.internal.com` or any other protected domain in the same realm. Cookie hijacking protection restricts cookies to the fully qualified domain name (FQDN) of the host where they are issued, such as `openam-server.example.com` and `server-with-agent.example.com`, using CDSSO to handle authentication and authorization.

For CDSSO with cookie hijacking protection, when a client successfully authenticates, AM issues the master SSO token cookie for its FQDN. AM issues *restricted token* cookies for the other FQDNs where the web or Java agents reside. The client ends up with cookies having different session identifiers for different FQDNs, and the AM server stores the correlation between the master SSO token and restricted tokens, such that the client only has one master session internally in AM.

To protect against cookie hijacking, you restrict the AM server domain to the server where AM runs. This sets the domain of the SSO token cookie to the host running the AM server that issued the token. You also enable use of a unique SSO token cookie. For your Java agents, you enable use of the unique SSO token cookie in the agent configuration.

Important

Client-based sessions do not support restricted tokens. Therefore, Web Agents and Java Agents configured in a realm configured for client-based sessions are not protected against cookie hijacking. ForgeRock recommends using web or Java agents with CTS-based sessions.

To Enable Restricted Tokens

1. In the AM console, navigate to Configure > Global Services > Platform.
 - a. Remove all domains from the Cookies Domains list.
 - b. Save your work.
2. Navigate to Configure > Server Defaults > Advanced.
3. Set the `com.sun.identity.enableUniqueSSOTokenCookie` advanced property to `true`.
4. Save your work.
5. Restart AM or the container in which it runs for the configuration changes to take effect.

Chapter 7

Additional Cookie Security Considerations

Although the session cookie is the most important cookie to keep track of when securing AM, there are other points you must consider, such as:

- Which cookie are you using for sticky load balancing?

By default, AM creates the `amlbcookie` cookie and sets it to the ID of the instance that first responded to a request. You should change the name of this cookie to something unique in your environment.

- Which other cookies, relevant for your environment, interact with AM or are sent to AM as part of a chain of requests?

The following table summarizes the tasks and information you need to review to manage cookie security that is not strictly related to the session cookie:

Task	Resources
Enable Support for <code>SameSite</code> Rules Configure AM to apply <code>SameSite</code> rules, such that you can declare that your cookies are restricted to a first-party or a same-site context.	"Enabling SameSite Cookie Rules"
Review the Secure Cookie Filter AM provides a filter that upgrades cookies to secure cookies if the conditions are met.	"Managing the Secure Cookie Filter"
Change the Name of the Sticky Load Balancing Cookie Name the cookie something relevant and unique for your environment.	"Changing the Name of the Sticky Load Balancing Cookie"

Enabling SameSite Cookie Rules

For additional cookie security, enable support for applying *SameSite* cookie rules, as described in the internet-draft [Cookies: HTTP State Management Mechanism](#).

You can configure the AM server to apply `SameSite` cookie rules by navigating to **Configure > Server Defaults > Advanced**, and setting the `com.sun.identity.cookie.samesite` property's value to one of the following:

strict

Requests originating from different sites will not have cookies sent with them.

When this mode is enabled, any AM functionality that relies on requests being redirected back to the AM instance may not operate correctly. For example, OAuth 2.0 flows and SAML federation may not operate correctly if AM cannot access the required cookies.

lax

Cookies received from different sites cannot be accessed, unless the request is using a *top-level* request, and uses a "safe" HTTP method, such as GET, HEAD, OPTIONS, or TRACE.

off

No restrictions on the domain of cookies is applied. This is the default setting.

You *must* disable `SameSite` support if any of the following is true:

- You have set `Access-Control-Allow-Credentials=true` in your CORS configuration. For more information on configuring CORS in AM, see "Configuring CORS Support".
- You are using SAML HTTP-POST bindings. For example, IDP-initiated single logout (SLO) functionality will not operate correctly if `SameSite` support is enabled, as the `iPlanetDirectoryPro` cookie would not be accessible in cross-domain POST requests. For more information on SAML single logout, see "Implementing SSO and SLO" in the *SAML v2.0 Guide*.

Caution

Modern browsers only allow disabling `SameSite` if the cookie is marked as `Secure`. If you need to handle cross-site requests with cookies, you should move to HTTPS-only environment.

Managing the Secure Cookie Filter

As part of the support that AM provides for `SameSite` cookies, the deployment descriptor file `web.xml` includes a filter that flags cookies as secure if any of the following is true:

- The request comes in through a connection marked as secure. For example, because you have marked an HTTP connector as secure in Tomcat.
- The request comes in through an HTTPS connector.

Automatically promoting cookies to secure ensures that the functionality continues to work with the `SameSite` changes, because you can only opt out of `SameSite` if a cookie is marked as secure.

To Exclude Cookies From the Filter

1. To exclude cookies from the filter, edit the `/path/to/tomcat/webapps/openam/WEB-INF/web.xml` file and search for the `SecureCookieFilter` filter.

2. Add any cookies you want to exclude to the list. For example:

```
...
<param-name>excludes</param-name>
<param-value>
  myCookie1
  myStickyCookie
  myCookie2
</param-value>
...
```

Tip

To ensure that non-secure requests are load-balanced correctly, the `amlbcookie` cookie is already excluded by default. If you are using a custom cookie for sticky load balancing, you may want to add it to the list of excluded cookies.

3. Restart AM or the container where it runs for the changes to take effect.

Changing the Name of the Sticky Load Balancing Cookie

By default, the sticky load balancing cookie name is `amlbcookie`. Change this value to something that is unique in your environment, and configure the name of the cookie in your load balancers to achieve session stickiness.

Perform the following steps to change the name of the cookie:

1. Go to Configure > Server Defaults > Advanced.
2. Change the value of the `com.ipplanet.am.lbcookie.name` advanced server property to the new cookie name.

By default, AM sets the value of the load balancing cookie to the ID of the instance that first responded to a request. You can change it, but we recommend that you keep this configuration when using Web Agents.

For more information, see "To Configure Site Sticky Load Balancing" in the *Setup Guide*.

3. Restart AM or the container where it runs.

Chapter 8

Securing Sessions

Cookie hijacking is not the only danger to sessions. Consider the following non-exhaustive list of scenarios that can result in a compromised account:

- End users entering their data in a malicious website thinking it is the authentic one.
- End users leaving their computers unattended while their session is open.
- End users logging in from completely different locations or devices than their usual.

The following table summarizes the tasks you need to perform to keep sessions secure:

Task	Resources
<p>Configure Settings Related to Session Termination</p> <p>Understand how session termination works in AM, and configure the session time-to-live and idle timeout.</p> <p>Ensuring sessions expire within a reasonable time helps you protect your environment against impersonation attacks.</p>	"Understanding Session Termination"
<p>Lock Users After Failed Login Attempts</p> <p>Configure account lockout to protect your environment against brute-force or dictionary attacks.</p>	"Configuring Account Lockout"
<p>Limit the Number of Active Sessions for a User</p> <p>Prevent users from logging in from more than two devices as a time, for example. This helps you mitigate against cases where user accounts have been compromised.</p>	"Configuring Session Quotas"
<p>Protect Client-Based Sessions</p> <p>AM offers additional security measures to protect client-based sessions. They are more vulnerable to hijacking than CTS-based sessions because they contain all the session information in them.</p>	"Configuring Client-Based Session Security"
<p>Protect Authentication Sessions</p> <p>Configure authentication session whitelisting to protect these sessions against replay attacks.</p>	"Configuring Client-Based Session Security"

Task	Resources
Protect Sensitive Attributes (Self-Service) Prevent attackers from changing sensitive attributes if they do hijack a session.	"Protect Sensitive Attributes"

Understanding Session Termination

AM manages active sessions, allowing single sign-on when authenticated users attempt to access system resources in AM's control.

AM ensures that user sessions are terminated when a configured timeout is reached, or when AM users perform actions that cause session termination. Session termination effectively logs the user out of all systems protected by AM.

With CTS-based sessions, AM terminates sessions in four situations:

- When a user explicitly logs out.
- When an administrator monitoring sessions explicitly terminates a session.
- When a session exceeds the maximum time-to-live.
- When a user is idle for longer than the maximum session idle time.

Under these circumstances, AM responds by removing CTS-based sessions from the CTS token store and from AM server memory caches. With the user's session no longer present in CTS, AM forces the user to reauthenticate during subsequent attempts to access resources protected by AM.

When a user explicitly logs out of AM, AM also attempts to invalidate the `iPlanetDirectoryPro` cookie in users' browsers by sending a `Set-Cookie` header with an invalid session ID and a cookie expiration time that is in the past. In the case of administrator session termination and session timeout, AM cannot invalidate the `iPlanetDirectoryPro` cookie until the next time the user accesses AM.

Session termination differs for client-based sessions. Since client-based sessions are not maintained in the CTS token store, administrators cannot monitor or terminate them. Because AM does not modify the `iPlanetDirectoryPro` cookie for client-based sessions after authentication, the session idle time is not maintained in the cookie. Therefore, AM does not automatically terminate client-based sessions that have exceeded the idle timeout.

As with CTS-based sessions, AM attempts to invalidate the `iPlanetDirectoryPro` cookie from a user's browser when the user logs out. When the maximum session time is exceeded, AM also attempts to invalidate the `iPlanetDirectoryPro` cookie in the user's browser the next time the user accesses AM.

It is important to understand that AM cannot guarantee cookie invalidation. For example, the HTTP response containing the `Set-Cookie` header might be lost. This is not an issue for CTS-based sessions, because a logged out session no longer exists in the CTS token store, and a user who attempts to access AM after previously logging out will be forced to reauthenticate.

However, the lack of a guarantee of cookie invalidation is an issue for deployments with client-based sessions. It could be possible for a logged out user to have an `iPlanetDirectoryPro` cookie. AM could not determine that the user previously logged out. Therefore, AM supports a feature that takes additional action when users log out of client-based sessions. AM can maintain a list of logged out client-based sessions in a session blacklist in the CTS token store. Whenever users attempt to access AM with client-based sessions, AM checks the session blacklist to validate that the user has not, in fact, logged out.

Since AM does not modify client-based session cookies once they are stored in the end user's browser, and client-based sessions contain, among others, the session maximum time-to-live, it is imperative to protect them against tampering. See "Configuring Client-Based Session Security" for more information.

Configuring Maximum Session Time-to-Live

When configuring the maximum session time-to-live, you must balance security and user experience. Depending on your application, it may be acceptable for your users to log in once a month. Financial applications, for example, tend to expire their sessions in less than an hour.

The longer a session is valid, the larger the window during which a malicious user could impersonate a user if they were able to hijack a session cookie.

To Configure Session Maximum Time-to-Live

1. Log in to the AM console as an administrative user, for example, `amAdmin`.
2. Navigate to Realms > Realm Name > Services > Session > Dynamic Attributes.

Note that you can also change maximum session time settings globally for the AM site by navigating to Configure > Sessions > Dynamic Attributes.

3. On the Maximum Session Time property, configure a value suitable for your environment.
4. Save your changes.

Configuring CTS-Based Session Idle Timeout

Consider a user with a valid session navigating through pages or making changes to the configuration. If for any reason they leave their desk and their computer remains open, a malicious user could take the opportunity to impersonate them.

Session idle timeout can help mitigate those situations, by logging out users after a specified duration of inactivity. Session idle timeout can only be used in realms configured for CTS-based sessions.

To Configure Session Idle Timeout

1. Log in to the AM console as an administrative user, for example, `amAdmin`.

2. Navigate to Realms > Realm Name > Services > Session > Dynamic Attributes.

Note that you can also change idle timeout settings globally for the AM site by navigating to Configure > Sessions > Dynamic Attributes.

3. On the Maximum Time Idle property, configure a value suitable for your environment.
4. Save your changes.

Configuring Client-Based Session Blacklisting

Session blacklisting ensures that users who have logged out of client-based sessions cannot achieve single sign-on without reauthenticating to AM. Session blacklisting does not apply to authentication sessions.

To Configure Session Blacklisting

1. Make sure that you deployed the Core Token Service during AM installation. The session blacklist is stored in the Core Token Service's token store.
2. Navigate to Configure > Global Services, click Session, and then locate the Client-based Sessions tab.
3. Select the Enable Session Blacklisting option to enable session blacklisting for client-based sessions. When you configure one or more AM realms for client-based sessions, you should enable session blacklisting in order to track session logouts across multiple AM servers.

Changing the value of this property takes effect immediately.

4. (Optional) Configure the Session Blacklist Cache Size property.

AM maintains a cache of logged out client-based sessions. The cache size should be around the number of logouts expected in the maximum session time. Change the default value of 10,000 when the expected number of logouts during the maximum session time is an order of magnitude greater than 10,000. An underconfigured session blacklist cache causes AM to read blacklist entries from the Core Token Service store instead of obtaining them from cache, which results in a small performance degradation.

Changing the value of this property takes effect immediately.

5. Configure the Blacklist Poll Interval property.

AM polls the Core Token Service for changes to logged out sessions if session blacklisting is enabled. By default, the polling interval is 60 seconds. The longer the polling interval, the more time a malicious user has to connect to other AM servers in a cluster and make use of a stolen session cookie. Shortening the polling interval improves the security for logged out sessions, but might incur a minimal decrease in overall AM performance due to increased network activity.

Changing the value of this property does not take effect until you restart AM.

6. Configure the Blacklist Purge Delay property.

When session blacklisting is enabled, AM tracks each logged out session for the maximum session time plus the blacklist purge delay. For example, if a session has a maximum time of 120 minutes and the blacklist purge delay is one minute, then AM tracks the session for 121 minutes. Increase the blacklist purge delay if you expect system clock skews in a cluster of AM servers to be greater than one minute. There is no need to increase the blacklist purge delay for servers running a clock synchronization protocol, such as Network Time Protocol.

Changing the value of this property does not take effect until you restart AM.

7. Click Save Changes.

Important

Enabling or disabling the session blacklist, or changing the cache size takes effect immediately.

Changes to any other session blacklist properties **do not** take effect until you restart AM.

For detailed information about Session Service configuration attributes, see the entries for "Session" in the *Reference*.

Configuring Account Lockout

Account lockout is a security mechanism that locks a user after repeated failed login attempts. It is used to slow down brute-force attacks as well as to compensate for weak password policies.

Most deployments use the identity store's password policy to control account lockout. If that is not an option to your deployment, configure account lockout in AM as explained in this section.

AM supports two different approaches to *account lockout*, where AM locks an account after repeated authentication failures; persistent lockout and memory lockout:

- Persistent (physical) lockout sets the user account status to **inactive** in the user profile. For persistent lockout, AM tracks failed authentication attempts by writing to the user repository.

Persistent account lockout works independently of account lockout mechanisms in the underlying directory server that serves as the user data store.

- Memory lockout locks the user account, keeping track of the locked state only in memory, and then unlocking the account after a specified delay. Memory lockout is also released when AM restarts.

Note

Failed login attempts during the transactional authorization flow do not increment account lockout counters. For more information on transactional authorization, see "*Transactional Authorization*" in the *Authorization Guide*.

If login failures are stored in AM's memory, this may result in user accounts not being locked out even after multiple login failures. To avoid this issue, make sure to implement persistent lockout instead.

To Configure Account Lockout

- To configure account lockout, configure the authentication settings for the realm, as follows:
 - a. Access the settings in the AM console under Realms > *Realm Name* > Authentication > Settings > Account Lockout.
 - b. Enable lockout by checking Login Failure Lockout Mode, setting the number of attempts, and setting the lockout interval and duration.

You can also opt to warn users after several consecutive failures, or to multiply the lockout duration on each successive lockout.
 - c. (Optional) If you have configured CTS-based or client-based authentication sessions, ensure the Store Invalid Attempts in Data Store switch is enabled. Failure to do so may result in users not being locked out even after multiple login failures.
 - d. To save account login failures to the Data Store, enable Store Invalid Attempts in Data Store. This step is necessary when using CTS-based or client-based authentication sessions.
 - e. You can set up email notification upon lockout to an administrator if AM is configured to send mail. You can configure AM to send mail in Configure > Server Defaults > General > Mail Server.
 - f. For persistent lockout, AM sets the value of the user's `inetuserstatus` profile attribute to `inactive`. You can also specify another attribute to update on lockout. You can further set a non-default attribute on which to store the number of failed authentication attempts. When you do store the number of failed attempts in the data store, other AM servers accessing the user data store can also see the number.

For more information, see "Configuring Realm Authentication Properties" in the *Authentication and Single Sign-On Guide*.

Tip

To unlock a user's account, find the user under Realms > *Realm Name* > Identities. Select the user you want to unlock, and set their User Status property to Active. Save your changes.

For information on how authentication trees handle account lockout, see "About Account Lockout for Trees" in the *Authentication and Single Sign-On Guide*.

Configuring Session Quotas

AM lets you limit the number of active sessions for a user by setting session quotas. Use this feature, for example, to prevent a user from logging in from more than two devices at once, mitigating scenarios where user passwords may have been compromised.

AM's support for session quotas requires CTS-based sessions.

To Configure Session Quotas and Exhaustion Actions

The session quota applies to all sessions opened for the same user (as represented by the user's universal identifier). To configure session quotas and exhaustion in AM, perform the following steps:

1. Log in to the AM console as an administrative user, for example, `amAdmin`.
2. Navigate to Configure > Global Services > Sessions > Session Quotas.
3. From the Enable Quota Constraints drop-down menu, select **ON**.
4. On the Set Resulting behavior if session quota exhausted property, set one of the following values:

DENY_ACCESS

Deny access, preventing the user from creating an additional session.

DESTROY_NEXT_EXPIRING

Remove the next session to expire, and create a new session for the user. The next session to expire is the session with the minimum time left until expiration.

This is the default setting.

DESTROY_OLDEST_SESSION

Remove the oldest session, and create a new session for the user.

DESTROY_OLD_SESSIONS

Remove all existing sessions, and create a new session for the user.

If none of these session quota exhaustion actions fit your deployment, you can implement a custom session quota exhaustion action. For an example, see "Customizing CTS-Based Session Quota Exhaustion Actions".

5. Navigate to Realms > *Realm Name* > Services > Session.
6. On the Set Active User Sessions property, configure the maximum number of concurrent sessions a user can have.

Note that you can also change this setting globally for the AM site in **Configure > Sessions > Dynamic Attributes**.

7. Save your work.

Configuring Client-Based Session Security

AM can issue CTS-based sessions, which contain a reference to the real session stored in the CTS store, or client-based sessions, which contain all the information that would be held in the CTS store.

While both types are susceptible to cookie hijacking, client-based sessions are even more vulnerable, since they contain all the information for the session. Therefore, the malicious user could tamper with the session data to their benefit.

When using client-based sessions and client-based authentication sessions, you should configure AM to sign and/or encrypt the JWT containing session information:

JWT Signing

AM verifies that the JWT is authentic by validating a signature configured in the Session Service. AM thwarts attackers who might attempt to tamper with the contents of the JWT or its signature, or who might attempt to sign the JWT with an incorrect signature.

JWT Encryption

Knowledgeable users can easily decode JWTs. Because an AM session or authentication session contains information that might be considered sensitive, encrypting the JWT that contains it protects its contents, ensuring opaqueness.

Encrypting the JWT prevents man-in-the-middle attacks that could log the state of every AM session. Encryption also ensures that end users are unable to access the information in their AM session.

Signing and encryption require certificates and shared secrets. When deploying multiple AM servers in a site, every server must have the same security configuration. Shared secrets and security keys must be identical. If you modify shared secrets or keys, you must update all the servers in the site. For more information about managing certificates for AM, see "*Configuring Secrets, Certificates, and Keys*".

Client-based sessions and authentication sessions share the same encryption and signing configuration for the realm.

Important

To ensure that the client-based session cookie size does not surpass the browser supported size, Web Agents and Java Agents do not support both signing and encrypting the session cookie. For more information, see "Client-Based Session Security and Agents".

The following sections provide detailed steps for configuring client-based session and authentication session cookie security:

- "Configuring the JWT Signature"
- "Configuring JWT Encryption"
- "Client-Based Session Security and Agents"

Configuring the JWT Signature

Configure a JWT signature to prevent malicious tampering of client-based session and authentication session JWTs.

Perform the following steps to configure the JWT signature:

To Configure the JWT Signature

1. Navigate to Configure > Global Services > Session > Client-based Sessions.
2. From the Signing Algorithm Type drop-down menu, select a suitable algorithm for your environment.

The default value is **RS256**.

Note

You should disable signing the JWT if you plan to encrypt the JWT with the Direct AES Encryption algorithm. To disable JWT signing, perform the following steps:

1. Navigate to Deployment > Servers > *Server Name* > Advanced.
2. Set the `org.forgerock.openam.session.stateless.signing.allownone` property to **true**.
3. Save your changes.
4. Navigate to Configure > Global Services > Session > Client-based Sessions.
5. From the Signing Algorithm Type drop-down list, select **NONE**.

6. Save your changes.

3. If you specified an HMAC signing algorithm, change the value in the Signing HMAC Shared Secret field if you do not want to use the generated default value.
4. If you specified the RS256 signing algorithm, specify a value in the Encryption RSA Certificate Alias field to use for signing the JWT signature.
5. Save your changes.

For detailed information about Session Service configuration attributes, see the entries for "Session" in the *Reference*.

Configuring JWT Encryption

Configure JWT encryption to prevent man-in-the-middle attackers from accessing users' session details, and to prevent end users from examining the content in the JWT.

Perform the following steps to encrypt the JWT:

To Configure JWT Encryption

1. Navigate to Configure > Global Services > Session > Client-based Sessions.
2. From the Encryption Algorithm drop-down list, select a suitable algorithm.
3. If you selected **RSA** in the previous step, you can select one of three padding options using the advanced property `org.forgerock.openam.session.stateless.rsa.padding`:
 - **RSA1_5**. RSA with PKCS#1 v1.5 padding.
 - **RSA-OAEP**. RSA with OAEP and SHA-1.
 - **RSA-OAEP-256**. RSA with OAEP padding and SHA-256.
 - a. In the AM console, select Configure > Server Defaults > Advanced.
 - b. In the Add a Value field, enter: `org.forgerock.openam.session.stateless.rsa.padding`.
 - c. In the corresponding Add a Value field, enter one of the padding options. For example, **RSA-OAEP**. The default is **RSA-OAEP-256**.
 - d. Click the plus sign ("+"), and then click Save Changes.
4. For the underlying content encryption method, select one of the following encryption methods supported in AM:
 - **A128CBC-HS256**. AES 128-bit in CBC mode with HMAC-SHA-256-128 hash (HS256 truncated to 128 bits)

- **A192CBC-HS384**. AES 192-bit in CBC mode with HMAC-SHA-384-192 hash (HS384 truncated to 192 bits)
- **A256CBC-HS512**. AES 256-bit in CBC mode with HMAC-SHA-512-256 hash (HS512 truncated to 256 bits)
- **A128GCM**. AES 128-bit in GCM mode
- **A192GCM**. AES 192-bit in GCM mode
- **A256GCM**. AES 256-bit in GCM mode

- a. In the AM console, select Configure > Server Defaults > Advanced.
 - b. In the Add a Value field, enter: `org.forgerock.openam.session.stateless.encryption.method`.
 - c. In the corresponding Add a Value field, enter one of the padding options. For example, `A128CBC-HS512`. The default is `A128CBC-HS256`.
 - d. Click the plus sign ("+"), and then click Save Changes.
5. In the Encryption RSA Certificate Alias field, enter an alias value to use for encrypting the JWT signature.

AM retrieves the certificate from the keystore specified by the `com.sun.identity.saml.xmlsig.keystore` property.

6. If you selected AES KeyWrapping or Direct AES Encryption, enter the key in the Encryption Symmetric AES Key field.

This should be a base64-encoded random key. For direct encryption with AES-GCM or for AES-KeyWrap with any content encryption method, this should be 128, 192 or 256 bits.

For direct encryption with AES-CBC-HMAC it should be double those sizes (one half for the AES key, the other half for the HMAC key).

7. Save your changes.
8. Ensure that the JWT signature configuration is identical on every AM server in your AM site.
9. (Optional) To compress the session state, select `Deflate Compression` from the Compression Algorithm drop-down list.

Warning

When set to `Deflate compression`, this option may lead to a possible vulnerability with session state information leakage. Because the session token compression depends on the data in the session, an attacker can vary one part of the session (for example, the username or some other property) and then deduce some secret parts of the session state by examining how the session compresses. You should evaluate this threat depending on your use cases before enabling compression and encryption together.

By default, AM rejects compressed session JWTs that expand to a size larger than 32 KiB (32768 bytes). For more information, see "Controlling the Maximum Size of Compressed JWTs".

For detailed information about Session Service configuration attributes, see the entries for "Session" in the *Reference*.

Client-Based Session Security and Agents

To ensure that the client-based session cookie size does not surpass the browser supported size, Web Agents and Java Agents do not support both signing and encrypting the session cookie. To configure agents with client-based sessions, implement one of the following configurations:

- Configure signing and compression:
 1. Enable HS256 signing for the client-based session cookie. For more information, see "[To Configure the JWT Signature](#)".
 2. Enable compression by navigating to Configure > Global Services > Session > Client-based Sessions and selecting **Deflate Compression** from the Compression Algorithm drop-down.
- Configure encryption and compression:
 1. Set the `org.forgerock.openam.session.stateless.signing.allownone` advanced server property to **true** on every AM instance in the site.
 2. Disable signing for the client-based session cookie by navigating to Configure > Global Services > Session > Client-based Sessions and selecting **NONE** from the Signing Algorithm Type drop-down.
 3. Enable Direct AES Encryption. For more information, see "[To Configure JWT Encryption](#)".
 4. Enable compression by navigating to Configure > Global Services > Session > Client-based Sessions and selecting **Deflate Compression** from the Compression Algorithm drop-down.

Failure to set up client-based sessions correctly may cause unexpected errors when accessing a protected resource, such as blank pages and redirection loops.

Client-based sessions do not support restricted tokens. Therefore, Web Agents and Java Agents configured in a realm configured for client-based sessions are not protected against cookie hijacking. ForgeRock recommends using web or Java agents with CTS-based sessions.

Protect Sensitive Attributes

If you configure user self-service, you *must* ensure that the user's email address and phone number cannot be changed without re-authentication. If you do not do this, an attacker that gains access to a user's session can change the user's email address and perform a password reset to gain full access to their account.

- To protect sensitive self-service attributes globally, select Configure > Services > Global Services > User Self Service > Profile Management and add **telephoneNumber** and **mail** to the list of Protected User Attributes.

- To protect sensitive self-service attributes at the realm level, select Realms > Realm name > Services > User Self Service > Profile Management and add `telephoneNumber` and `mail` to the list of Protected User Attributes.

For more information, see "Profile Management"

Configuring Authentication Session Whitelisting

Enable authentication session whitelisting to protect authentication sessions from replay attacks.

When authentication session whitelisting is enabled, AM generates a key-value pair for each authentication session and stores it for the length of the authentication flow in the following ways:

- For client-based authentication sessions, AM stores the key-value pair in the CTS token store.
- For CTS-based authentication sessions, AM creates the key-value pair as a session property in the authentication session.
- For in-memory sessions, AM creates the key-value pair as a session property in the authentication session.

Each time the authentication flow reaches an authentication node, AM modifies the value of the stored key-value pair and sends it to the user or client that it is authenticating. The next request to AM to continue the authentication flow must contain the key-value pair and must match the value expected by AM.

If the authenticating user or client cannot provide the key-value pair with the values AM expects, AM would not continue the authentication flow, therefore protecting the authentication flow against malicious users wanting to rewind the authentication flow to a previous node.

Perform the following steps to configure authentication session whitelisting:

To Configure Authentication Session Whitelisting

1. Navigate to Realms > *Realm Name* > Authentication > Settings > Trees.
2. Select Enable whitelisting.
3. Save your changes.

Chapter 9

Request Security Considerations

AM can receive requests from multiple sources and for different purposes, such as authentication requests, RESTful requests to the endpoints, and POSTs that potentially may include a lot of data.

Containers usually have settings to mitigate against DoS attacks that POST large amounts of form data to your applications. Refer to your container documentation for more information about their settings, and how they can protect AM.

These settings, however, do not protect AM from receiving large amounts of POST data from other sources.

The following table summarizes the steps AM takes to protect against being overloaded, and how to adjust default values:

Task	Resources
Controlling the Maximum Size of Decompressed JWTs By default, AM rejects JWTs that expand to a size larger than 32 KiB (32768 bytes) when decrypted.	"Controlling the Maximum Size of Compressed JWTs".
Limiting the Size of the Request Body By default, AM rejects incoming requests whose body is larger than 1 MB (1048576 bytes) in size.	"Limiting the Size of the Request Body".

Controlling the Maximum Size of Compressed JWTs

A number of AM features accept JWTs to receive information. Some examples are:

- "*The Remote Consent Service*" in the *OAuth 2.0 Guide*, when it receives consent responses.
- The OAuth 2.0/OpenID Connect authorization service, when:
 - OpenID Connect clients send **request** in the *OAuth 2.0 Guide* parameters as a JWT instead of as HTTP parameters.
 - OpenID Connect clients register dynamically using **software statements** in the *OpenID Connect 1.0 Guide*.
- The Authentication service, when configured to issue **client-based sessions** in the *Sessions Guide*.

These JWTs that AM receives can be signed and/or encrypted. Sometimes, larger JWTs are compressed to improve delivery speeds to AM.

Decompressing a JWT makes it expand in size. By default, AM rejects any JWT that expands to more than 32 KiB (32768 bytes), and throws an exception with a message similar to `JWT payload decompressed to larger than maximum allowed size`.

Ensure that the JWTs your clients send to AM are smaller than 32 KiB before compression, or increase the 32 KiB value to a reasonable limit. Take into account that AM performs decryption and decompression operations in its heap, and that you do not want to allow very large JWTs to, potentially, leave AM out of memory.

If you need to change the default value, perform the following steps:

1. Configure the `org.forgerock.json.jose.jwe.compression.max.decompressed.size.bytes` Java system property on the container where AM runs.

For example, edit the `setenv.sh` file of the Apache Tomcat instance, and set the property with the new size in bytes:

```
JAVA_OPTS="$JAVA_OPTS -Dorg.forgerock.json.jose.jwe.compression.max.decompressed.size.bytes=40960"
```

2. Restart the container for the changes to make effect.

Limiting the Size of the Request Body

HTTP requests are not limited by the specification. Rather, the method used limits the amount of data that a client can send. The GET and DELETE methods, for example, are limited by the size of the URL. The POST method is not. Instead, browsers and application servers limit the amount of data a request can send to your applications.

Ensure that the amount of data that reaches your applications and AM is not large enough to overwhelm them.

Application servers usually can mitigate against DoS attacks that POST large amounts of form data, but AM endpoints may receive large amounts of POST data in different ways, such as in JSON, JWT, or JWK formats.

By default, AM rejects incoming requests with a body larger than 1 MB (1048576 bytes) in size. It also returns an HTTP 413 error response, and logs a message similar to the following:

- `ERROR: Request Content-Length exceeds maximum allowed`, if the content's length was specified in the request.
- `ERROR: Counted request entity size exceeds maximum allowed`, if the content's length was not specified.

To change the default value, perform the following steps:

- Change the value of the `org.forgerock.openam.request.max.bytes.entity.size` advanced server property to the new size, in bytes.

+ *How Do I Configure Advanced Server Properties?*

- To configure advanced server properties in the AM Admin UI for all AM instances, go to Configure > Server Defaults > Advanced.
- To configure advanced server properties for a particular instance, go to Deployment > Servers > *Server Name* > Advanced.

If the property you want to add or edit is not already configured, add it with its value, then click on the plus (+) button.

If the property you want to add or edit is already configured, click on the pencil (✎) button to edit it. When you are finished, click on the tick (✓) button.

Save your changes.

The property is hot-swappable. You do not need to restart AM for the changes to take effect.

Chapter 10

Protecting Applications

AM provides authentication and authorization capabilities, but it requires a policy enforcement point (PEP) intercepting traffic to the applications.

ForgeRock offers Java Agents, Web Agents, and IG as PEPs to enforce what AM decides in a way that is unobtrusive to the user.

Identity Gateway or AM Web and Java Agents?

ForgeRock Identity Gateway and the AM web and Java agents can both enforce policy, redirecting users to authenticate when necessary, and controlling access to protected resources. IG runs as a self-contained reverse proxy located between the users and the protected applications. Web and Java agents are installed into the servers where applications run, intercepting requests in that context.

Use IG to protect access to applications not suited for a web or Java agent, for example, those applications deployed on operating systems or web servers or containers not supported by the agents.

Web and Java agents have the advantage of sitting within your existing server infrastructure. Once you have agents installed into the servers with web applications or sites to protect, then you can manage their configurations centrally from AM.

For organizations with both servers on which you can install web and Java agents and also applications that you must protect without touching the server, you can use agents on the former and IG for the latter.

For more information about Agents, see the *ForgeRock Web Agents User Guide*, or the *ForgeRock Java Agents User Guide*.

For more information about IG, see the *ForgeRock Identity Gateway* documentation.

Chapter 11

Setting Up Audit Logging

AM supports a comprehensive Audit Logging Service that captures key auditing events, critical for system security, troubleshooting, and regulatory compliance.

Audit logs gather operational information about events occurring within an AM deployment to track processes and security data, such as authentication mechanisms, system access, user and administrator activity, error messages, and configuration changes.

This chapter describes the ForgeRock® Common REST-based Audit Logging Service available in AM. AM also supports a legacy Logging Service, based on a Java SDK. The legacy Logging Service will be deprecated in a future release of AM.

The Audit Logging Service uses a structured message format that adheres to a consistent log structure common across the ForgeRock Identity Platform. This common structure allows correlation between log messages of the different components of the Platform once the transaction IDs are trusted by enabling the ForgeRock trust transaction header system property.

For more information, see "Configuring the Trust Transaction Header System Property".

Important

The DS JSON logger is enabled by default. However, the ForgeRock transaction IDs are not trusted initially. You must set `trust-transaction-ids:true` to correlate log messages in DS with log messages in AM. For more information, see *To Configure LDAP JSON Access Logs* in the *ForgeRock Directory Services Monitoring Guide*.

Review the following sections to understand how audit logging works in AM, and how to implement it:

Task	Resources
Discover AM's Audit Logging Service AM auditing service provides a rich set of features to help you capture events that are critical for system security, troubleshooting, and regulatory compliance.	See "About the Audit Logging Service".
Configure AM to Log Audit Events Decide how to implement your audit login service, either globally or by realm, and configure audit login handlers to store audit events into files, databases, or other stores.	See "Implementing the Audit Logging Service".
Audit Log Reference	See "Audit Logging Reference".

Task	Resources
Check the format of the files, the names of the events, and more.	

Note

AM also supports the classic Logging Service, which is deprecated. For more information, see "Implementing the Classic Logging Service".

About the Audit Logging Service

AM writes log messages generated from audit events triggered by its instances, web or Java agents, the **ssoadm** tool, and connected ForgeRock Identity Platform implementations.

AM's Audit Logging Service provides a versatile and rich feature set as follows:

- **Global and Realm-Based Log Configuration.** You can configure audit logging globally, which ensures that all realms inherit your global log settings. You can also configure audit logging by realm, which allows you to set different log settings for each realm.
- **Audit Event Handlers.** The Audit Logging Service supports a variety of audit event handlers that allow you to write logs to different types of data stores. See "Configuring Audit Event Handlers" for a list of event handlers available in AM.
- **Audit Event Buffering.** By default, AM writes each log message separately as they are generated. AM supports message buffering, a type of batch processing, that stores log messages in memory and flushes the buffer after a preconfigured time interval or after a certain number of log messages reaches the configured threshold value.
- **Tamper-Evident Logging** for the CSV audit event handler. You can digitally sign your audit to enable the detection of tampering.
- **Log Rotation and Retention Policies.** AM rotates JSON and CSV audit logs when it reaches a specified maximum size. You can also configure a time-based rotation policy, which disables the max-size rotation policy and implements log rotation based on a preconfigured time sequence. AM also provides the option to disable log rotation completely for these file types. AM does not support external log rotation for JSON and CSV audit logs.

For Syslog, JDBC, JMS, and Splunk handlers, AM does not control log rotation and retention as they are handled by each respective service.

- **Blacklist and Whitelist Support.** The Audit Logging Service supports whitelist and blacklist-filtering to show or hide sensitive values or fields in the logs, such as HTTP headers, query parameters, cookies, profile attributes, or the entire field value.
- **Reverse DNS Lookup.** The Audit Logging Service supports a reverse DNS lookup feature for network troubleshooting purposes. Reverse DNS lookup is disabled by default as it enacts a performance hit in operation throughput.

Audit Log Topics

AM integrates log messages based on four different audit topics. A *topic* is a category of audit log event that has an associated one-to-one mapping to a schema type. Topics can be broadly categorized as access details, system activity, authentication operations, and configuration changes. The following table shows the basic event topics and associated audit log files for AM's default audit logging configuration, which uses a JSON audit event handler:

Audit Log Topics

Event Topic	File Name	Description
Access	<code>access.audit.json</code>	Captures who, what, when, and output for every access request.
Activity	<code>activity.audit.json</code>	Captures state changes to objects that have been created, updated, or deleted by end users (that is, non-administrators). Session, user profile, and device profile changes are captured in the logs.
Authentication	<code>authentication.audit.json</code>	Captures when and how a subject is authenticated and related events.
Configuration	<code>config.audit.json</code>	Captures configuration changes to the product with a timestamp and by whom. Note that the <code>userId</code> indicating the subject who made the configuration change is not captured in the <code>config.audit.json</code> but may be tracked using the <code>transactionId</code> in the <code>access.audit.json</code> .

Audit Logging in Web and Java Agents

Web and Java agents log audit events for security, troubleshooting, and regulatory compliance. You can store web or Java agent audit event logs in the following ways:

- **Remotely.** Log audit events to the audit event handler configured in the AM realm.
- **Locally.** Log audit events to a file in the web or Java agent installation directory.

For more information about audit logs in agents, see the *ForgeRock Web Agents User Guide* and the *ForgeRock Java Agents User Guide*.

Implementing the Audit Logging Service

When implementing the Audit Logging Service, decide whether you require specific audit systems per realm, or if a global configuration suits your deployment. Next, determine which event handlers suit your needs from those supported by AM. See the following sections for more information:

- To configure the Audit Logging Service, see "Configuring Audit Logging".
- To configure the audit event handlers, see "Configuring Audit Event Handlers".
- To configure the propagation of the transaction ID across the ForgeRock Identity Platform, see "Configuring the Trust Transaction Header System Property".

AM also supports the classic Logging Service, based on Java SDK, that will be deprecated in a future release. For more information, see "Implementing the Classic Logging Service".

Configuring Audit Logging

AM's default audit event handler is the JSON audit event handler, which comes configured and enabled for the global Audit Logging Service. The global configuration is used to control audit logging in realms that do not have the Audit Logging Service added to them. AM also supports configuring an Audit Logging Service on a per-realm basis.

The JSON audit event handler stores its JSON log files under `/path/to/openam/var/audit/`.

- To modify the global audit logging configuration, see "To Configure Global Audit Logging".
- To override the global audit logging configuration for a realm, see "To Configure Audit Logging in a Realm".

To Configure Global Audit Logging

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Navigate to Configure > Global Services > Audit Logging.
3. Configure the following options on the Global Attributes tab:
 - a. Activate Audit logging to start the audit logging feature.
 - b. In the Field whitelist filters and Field blacklist filters lists, enter any values to include (whitelist) or exclude (blacklist) from the audit event logs.

AM has a predefined whitelist built in that only records values that do not contain sensitive information. Use the filters to override the built in list, or to hide additional values that you do not want recorded.

Important

The Audit Logging Service lets you suppress the output of certain event types, because logging them may have an impact on performance. These event types are not logged by default, regardless of the configuration of the filter lists.

The filter lists will only apply to these event topics if logging is enabled for them.

For more information, see [org.forgerock.openam.audit.identity.activity.events.blacklist](#) in "Advanced Properties" in the *Reference*.

For information about the fields that appear in the default whitelist, see "Audit Log Default Whitelist".

To specify an additional field or value to be whitelisted, or blacklisted, add a value using a JSON pointer-like syntax that starts with the event topic ([access](#), [activity](#), [authentication](#), or [config](#)), followed by the field name, or the path to the value in the field.

The lists allow two types of filtering:

- Filter fields in events. You may be interested in capturing, or hiding, HTTP headers, query parameters, or potentially sensitive data like passwords in the access logs.

For example, you might want to filter out surnames by hiding the [sn](#) field from [activity](#) events. To do so, add the following pointers to the Field blacklist filters list:

```
/activity/before/sn  
/activity/after/sn
```

- Filter specific values in fields that store key-value pairs as JSON, such as the HTTP headers, query parameters, and cookies.

For example, to include the [Accept-Language](#) value in the [http.request.headers](#) field in [access](#) events, add the following pointer to the Field whitelist filters list:

```
/access/http/request/headers/accept-language
```

c. Save your changes.

For more information on configuring audit logging properties, see "Audit Logging" in the *Reference*.

4. On the Secondary Configurations tab, you can edit the configuration of the Global JSON Handler, or you can create new audit event handlers. For more information, see "Configuring Audit Event Handlers".

To Configure Audit Logging in a Realm

You can configure the Audit Logging Service for realms, allowing you to configure realm-specific log locations and handler types.

When the Audit Logging Service is added to a realm, it inherits the configuration defined under Configure > Global Services > Audit Logging > Realm Defaults. Properties configured explicitly in the realm-level service override the realm defaults.

To configure the Audit Logging Service in a realm, perform the following steps:

1. Navigate to Realms > *Realm Name* > Services.
2. Select Add a Service.
3. From the Choose a service type drop-down list, select Audit Logging.
4. Select Create.

The Audit Logging Service page appears. Configure the Audit Logging Service as follows:

5. Ensure audit logging is Enabled.
6. In the Field whitelist filters and Field blacklist filters lists, enter any values to include (whitelist) or exclude (blacklist) from the audit event logs.

AM has a predefined whitelist built in that only records values that do not contain sensitive information. Use the filters to override the built in list, or to hide additional values that you do not want recorded.

Important

The Audit Logging Service lets you suppress the output of certain event types, because logging them may have an impact on performance. These event types are not logged by default, regardless of the configuration of the filter lists.

The filter lists will only apply to these event topics if logging is enabled for them.

For more information, see [org.forgerock.openam.audit.identity.activity.events.blacklist](#) in "Advanced Properties" in the *Reference*.

For information about the fields that appear in the default whitelist, see "Audit Log Default Whitelist".

To specify an additional field or value to be whitelisted, or blacklisted, add a value using a JSON pointer-like syntax that starts with the event topic ([access](#), [activity](#), [authentication](#), or [config](#)), followed by the field name, or the path to the value in the field.

The lists allow two types of filtering:

- Filter fields in events. You may be interested in capturing, or hiding, HTTP headers, query parameters, or potentially sensitive data like passwords in the access logs.

For example, you might want to filter out surnames by hiding the [sn](#) field from *activity* events. To do so, add the following pointers to the Field blacklist filters list:

```
/activity/before/sn
/activity/after/sn
```

- Filter specific values in fields that store key-value pairs as JSON, such as the HTTP headers, query parameters, and cookies.

For example, to include the `Accept-Language` value in the `http.request.headers` field in `access` events, add the following pointer to the Field whitelist filters list:

```
/access/http/request/headers/accept-language
```

7. Save your changes.

For more information on configuring audit logging properties, see "Audit Logging" in the *Reference*.

8. On the Secondary Configurations tab, select Add a Secondary Configuration. Choose an event handler from the list.

For more information about supported event handlers and how to configure them, see "Configuring Audit Event Handlers".

Configuring Audit Event Handlers

AM supports the following types of audit event handlers:

Audit Event Handlers

Audit Event Handler Type	Publishes to	How to Configure
JSON	JSON files	"Configuring JSON Audit Event Handlers"
CSV	CSV files	"Configuring CSV Audit Event Handlers"
Syslog	The syslog daemon	"Configuring Syslog Audit Event Handlers"
JDBC	A relational database	"Implementing JDBC Audit Event Handlers"
JMS	JMS topics	"Implementing JMS Audit Event Handlers"
Splunk	A Splunk server	"Implementing Splunk Audit Event Handlers"

This section provides procedures for configuring each type of audit handler.

Configuring JSON Audit Event Handlers

The following procedure describes how to configure a JSON audit event handler:

To Configure a JSON Audit Event Handler

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to `Configure > Global Services > Audit Logging`.

Note that the JSON audit event handler is already configured in the global configuration. Select it to change its properties.

- To create the event handler in a realm, navigate to Realms > *Realm Name* > Services > Audit Logging.
- 3. On the Secondary Configurations tab, click Global JSON Handler or the Edit icon on the right if present. If no handler is present, click Add a Secondary Configuration, and select JSON.
- 4. Under the New JSON configuration page, enter a name for the event handler. For example, **JSON Audit Event Handler**.
- 5. (Optional) In the Rotation Times field, enter a time duration after midnight to trigger file rotation, in seconds. For example, you can provide a value of **3600** to trigger rotation at 1:00 AM. Negative durations are not supported.
- 6. Click Create.

After the JSON audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

- 7. On the General Handler Configuration tab, enable the event handler and configure the topics for your audit logs:
 - a. Select Enabled to activate the event handler, if disabled.
 - b. Choose the topics for your audit logs. For a description of each topic, see "Audit Log Topics".
 - c. Click Save Changes.
- 8. On the JSON Configuration tab, configure JSON options:
 - a. Override the default location of your logs if necessary, and save your changes. The default value is **%BASE_DIR%/%SERVER_URI%/Log/**.

Important

Make sure to configure a different log directory for each JSON audit event handler instance. If two instances are writing to the same file, it can interfere with log rotation and tamper-evident logs.

- b. Select ElasticSearch JSON Format Compatible to direct AM to generate JSON formats that are compatible with the ElasticSearch format.
- c. In the File Rotation Retention Check Interval field, edit the time interval (seconds) to check the time-base file rotation policies.
- d. Click Save Changes.

9. On the File Rotation tab, configure how files are rotated when they reach a specified file size or time interval:
 - a. Select Rotation Enabled to activate file rotation. If file rotation is disabled, AM ignores log rotation and appends to the same file.
 - b. In the Maximum File Size field, enter the maximum size of an audit file before rotation.
 - c. (Optional). In the File Rotation Prefix field, enter an arbitrary string that will be prefixed to every audit log to identify it. This parameter is used when time-based or size-based rotation is enabled.
 - d. In the File Rotation Suffix field, enter a timestamp suffix based on the Java SimpleDateFormat that will be added to every audit log. This parameter is used when time-based or size-based log rotation is enabled. The default value is `-yyyy.MM.dd-kk.mm.ss`.
 - e. In the Rotation Interval field, enter a time interval to trigger audit log file rotation in seconds. A negative or zero value disables this feature.
 - f. (Optional) In the Rotation Times field, enter a time duration after midnight to trigger file rotation, in seconds. For example, you can provide a value of `3600` to trigger rotation at 1:00 AM. Negative durations are not supported.
 - g. Click Save Changes.
10. On the File Retention tab, configure how long log files should be retained in your system:
 - a. In the Maximum Number of Historical Files field, enter a number for allowed backup audit files. A value of `-1` indicates an unlimited number of files and disables the pruning of old history files.
 - b. In the Maximum Disk Space field, enter the maximum amount of disk space that the audit files can use. A negative or zero value indicates that this policy is disabled.
 - c. In the Minimum Free Space Required field, enter the minimum amount of disk space required to store audit files. A negative or zero value indicates that this policy is disabled.
 - d. Click Save Changes.
11. On the Buffering tab, configure whether log events should be buffered in memory before they are written to the JSON file:
 - a. In the Batch Size field, enter the maximum number of audit log events that can be buffered.
 - b. In the Write interval field, enter the time interval in milliseconds at which buffered events are written to a file.
 - c. Click Save Changes.

Configuring CSV Audit Event Handlers

The following procedure describes how to configure a comma-separated values (CSV) audit event handler:

To Configure a CSV Audit Event Handler

Important

Due to the security concerns of opening CSV files with Excel, OpenOffice, and other spreadsheet programs, it is recommended that you open CSV files with alternative software, such as a text editor.

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to `Configure > Global Services > Audit Logging`.

Note that the CSV audit event handler is already configured in the global configuration. Select it to change its properties.

- To create the event handler in a realm, navigate to `Realms > Realm Name > Services > Audit Logging`.
3. On the Secondary Configurations tab, select `Add a Secondary Configuration`. Choose `CVS` from the list.

The `New CVS` page appears. Enter the basic configuration for the new handler by performing the following actions:

4. Enter a name for the event handler. For example, `CSV Audit Event Handler`.
5. (Optional) In the `Rotation Times` field, enter a time duration after midnight to trigger file rotation, in seconds. For example, you can provide a value of `3600` to trigger rotation at 1:00 AM. Negative durations are not supported.
6. Enable or disable the `Buffering` option.
7. Select `Create`.

After the CSV audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

8. On the `General Handler Configuration` tab, enable the event handler and configure the topics for your audit logs:
 - a. Select `Enabled` to activate the event handler, if disabled.

- b. Choose the topics for your audit logs. For a description of each topic, see "Audit Log Topics".
 - c. Save your changes.
9. On the CSV Configuration tab, override the default location of your logs if necessary, and save your changes. The default value is `%BASE_DIR%/SERVER_URI%/log/`.

Important

Configure a different log directory for each CVS audit event handler instance. If two instances are writing to the same file, it can interfere with log rotation and tamper-evident logs.

10. On the File Rotation tab, configure how files are rotated when they reach a specified file size or time interval:
 - a. Select Rotation Enabled to activate file rotation. If file rotation is disabled, AM ignores log rotation and appends to the same file.
 - b. In the Maximum File Size field, enter the maximum size of an audit file before rotation.
 - c. (Optional). In the File Rotation Prefix field, enter an arbitrary string that will be prefixed to every audit log to identify it. This parameter is used when time-based or size-based rotation is enabled.
 - d. In the File Rotation Suffix field, enter a timestamp suffix based on the Java SimpleDateFormat that will be added to every audit log. This parameter is used when time-based or size-based log rotation is enabled. The default value is `-yyyy.MM.dd-kk.mm.ss`.
 - e. In the Rotation Interval field, enter a time interval to trigger audit log file rotation in seconds. A negative or zero value disables this feature.
 - f. (Optional) In the Rotation Times field, enter a time duration after midnight to trigger file rotation, in seconds. For example, you can provide a value of `3600` to trigger rotation at 1:00 AM. Negative durations are not supported.
 - g. Save your changes.
11. On the File Retention tab, configure how long log files should be retained in your system:
 - a. In the Maximum Number of Historical Files field, enter a number for allowed backup audit files. A value of `-1` indicates an unlimited number of files and disables the pruning of old history files.
 - b. In the Maximum Disk Space field, enter the maximum amount of disk space that the audit files can use. A negative or zero value indicates that this policy is disabled.
 - c. In the Minimum Free Space Required field, enter the minimum amount of disk space required to store audit files. A negative or zero value indicates that this policy is disabled.

- d. Save your changes.
12. On the Buffering tab, configure whether log events should be buffered in memory before they are written to the CSV file:

- a. Select Buffering Enabled to activate buffering.

When buffering is enabled, all audit events are put into an in-memory buffer (one per handled topic), so that the original thread that generated the event can fulfill the requested operation, rather than wait for I/O to complete. A dedicated thread (one per handled topic) constantly pulls events from the buffer in batches and writes them to the CSV file. If the buffer becomes empty, the dedicated thread goes to sleep until a new item gets added. The default buffer size is 5000 bytes.

- b. Enable Flush Each Event Immediately to write all buffered events before flushing.

When the dedicated thread accesses the buffer, it copies the contents to an array to reduce contention, and then iterates through the array to write to the CSV file. The bytes written to the file can be buffered again in Java classes and the underlying operating system.

When Flush Each Event Immediately is enabled, AM flushes the bytes after each event is written. If the feature is disabled (default), the Java classes and underlying operation system determine when to flush the bytes.

- c. Save your changes.

13. On the Tamper Evident Configuration tab, configure whether to detect audit log tampering:

- a. Select Is Enabled to activate the tamper evident feature for CSV logs.

When tamper evident logging is enabled, AM generates an HMAC digest for each audit log event and inserts it into each audit log entry. The digest detects any addition or modification to an entry.

AM also supports another level of tamper evident security by periodically adding a signature entry to a new line in each CSV file. The entry signs the preceding block of events, so that verification can establish if any of these blocks have been added, removed, or edited by some user.

- b. In the Certificate Store Location field, enter the location of the keystore AM will use to sign the CSV logs, by default %BASE_DIR%/SERVER_URI%/Logger.jks.

The recommended approach is to create two keystores:

- A keystore for AM to use. This keystore is configured in the Certificate Store Location field and must contain a signing key pair called signature and an HMAC key called password.
- A keystore for the verification tool. This keystore must contain the HMAC password key, and the public key of the signature key pair.

You can use a simple script to create your keystores, for example: `create-keystore.sh`.

- c. In the Certificate Store Password field, enter the password of the keystore.
- d. In the Signature Interval field, enter a value in seconds for AM to generate and add a new signature to the audit log entry.
- e. Save your changes.
- f. To verify that rotated logs have not been tampered with, perform the following steps:
 - i. Download the AM-SSOAdminTools-5.1.3.11.zip file from the *ForgeRock BackStage* website.
 - ii. Install the tools by performing the steps in "To Set Up Administration Tools" in the *Installation Guide*.
 - iii. Use the **verifyarchive** tool to verify rotated log files as follows:

```
$ /path/to/AM-SSOAdminTools-5.1.3.11/openam/bin/verifyarchive \
--archive /path/to/openam/var/audit/ \
--topic access \
--suffix -yyyy.MM.dd-HH.mm.ss \
--keystore /path/to/keystore-verifier.jks \
--password password
```

In this example, the tool checks files with a suffix of the type `yyyy.MM.dd-HH.mm.ss` using their counterparts with suffix `.keystore`. For example, the tool checks the `tamper-evident-access-csv-2019.01.12-12.04.33` file against the `tamper-evident-access-csv-2019.01.12-12.04.33.keystore` file. The `.keystore` file contains the HMAC digest the tool uses to validate the signature of the logs.

The tool returns **PASS** or **FAIL** alongside the names of the files that have been tested. For example:

PASS	tamper-evident-access-csv-2019.01.12-12.04.33	
FAIL	tamper-evident-access-csv-2019.01.12-12.05.20	The HMac at row 2 is not correct.

Tip

Run the tool without any parameters to see the online help.

Configuring Syslog Audit Event Handlers

AM can publish audit events to a syslog server, which is based on a widely-used logging protocol. You can configure your syslog settings on the AM console.

The following procedure describes how to configure a Syslog audit event handler:

To Configure a Syslog Audit Event Handler

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to `Configure > Global Services > Audit Logging`.
 - To create the event handler in a realm, navigate to `Realms > Realm Name > Services > Audit Logging`.
3. On the Secondary Configurations tab, select `Add a Secondary Configuration`. Choose Syslog from the list.

The New Syslog page appears. Enter the basic configuration for the new handler by performing the following actions:

4. Enter a name for the event handler. For example, `Syslog Audit Event Handler`.
5. In the Server hostname field, enter the hostname or IP address of the receiving syslog server.
6. In the Server port field, enter the port of the receiving syslog server.
7. In the Connection timeout field, enter the number of seconds to connect to the syslog server. If the server has not responded in the specified time, a connection timeout occurs.
8. Enable or disable the Buffering option.
9. Select `Create`.

After the syslog audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

10. On the General Handler Configuration tab, enable the event handler and configure the topics for your audit logs:
 - a. Select `Enabled` to activate the event handler, if disabled.
 - b. Choose the topics for your audit logs. For a description of each topic, see ["Audit Log Topics"](#).
 - c. Save your changes.
11. On the Audit Event Handler Factory tab, keep the default class name for the audit event handler.
12. On the Syslog Configuration tab, configure the main syslog event handler properties:
 - a. In the Server hostname field, enter the hostname or IP address of the receiving syslog server.
 - b. In the Server port field, enter the port of the receiving syslog server.

- c. In the Connection timeout field, enter the number of seconds to connect to the syslog server. If the server has not responded in the specified time, a connection timeout occurs.
- d. From the Transport Protocol drop-down list, select TCP or UDP.
- e. Choose the facility.

A syslog message includes a PRI field that is calculated from the facility and severity values. All topics set the severity to **INFORMATIONAL** but you can choose the facility from the Facility drop-down list:

Syslog Facilities

Facility	Description
AUTH	Security or authorization messages
AUTHPRIV	Security or authorization messages
CLOCKD	Clock daemon
CRON	Scheduling daemon
DAEMON	System daemons
FTP	FTP daemon
KERN	Kernel messages
LOCAL0	Local use 0 (local0)
LOCAL1	Local use 1 (local1)
LOCAL2	Local use 2 (local2)
LOCAL3	Local use 3 (local3)
LOCAL4	Local use 4 (local4)
LOCAL5	Local use 5 (local5)
LOCAL6	Local use 6 (local6)
LOCAL7	Local use 7 (local7)
LOGALERT	Log alert
LOGAUDT	Log audit
LPR	Line printer subsystem
MAIL	Mail system
NEWS	Network news subsystem
NTP	Network time protocol
SYSLOG	Internal messages generated by syslogd
USER	User-level messages

Facility	Description
UUCP	Unix-to-unix-copy (UUCP) subsystem

- f. Save your changes.

13. On the Buffering tab, configure whether you want buffering or not:

- a. Select Buffering Enabled to activate it.

When buffering is enabled, all audit events that get generated are formatted as syslog messages and put into a queue. A dedicated thread constantly pulls events from the queue in batches and transmits them to the syslog server. If the queue becomes empty, the dedicated thread goes to sleep until a new item gets added. The default queue size is 5000.

- b. Save your changes.

Implementing JDBC Audit Event Handlers

You can configure AM to write audit logs to Oracle, MySQL, PostgreSQL, or other JDBC databases. AM writes audit log records to the following tables: `am_auditaccess`, `am_auditactivity`, `am_auditauthentication`, and `am_auditconfig`. For more information on the JDBC table formats for each of the logs, see "JDBC Audit Log Tables".

Before configuring the JDBC audit event handler, you must perform several steps to allow AM to log to the database:

To Prepare for JDBC Audit Logging

1. Create tables in the relational database in which you will write the audit logs. The SQL for Oracle, PostgreSQL, and MySQL table creation is in the `audit.sql` file under `/path/to/tomcat/webapps/openam/WEB-INF/template/sql/db-type`.

If you are using a different relational database, tailor one of the provided `audit.sql` files to conform to your database's SQL syntax.

2. JDBC audit logging requires a database user with read and write privileges for the audit tables. Do one of the following:
 - Identify an existing database user and grant that user privileges for the audit tables.
 - Create a new database user with read and write privileges for the audit tables.
3. Obtain the JDBC driver from your database vendor. Place the JDBC driver `.zip` or `.jar` file in the container's `WEB-INF/lib` classpath. For example, place the JDBC driver in `/path/to/tomcat/webapps/openam/WEB-INF/lib` if you use Apache Tomcat.

The following procedure describes how to configure a JDBC audit event handler. Perform the following steps after you have created audit log tables in your database and installed the JDBC driver in the AM web container:

To Configure a JDBC Audit Event Handler

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to `Configure > Global Services > Audit Logging`.
 - To create the event handler in a realm, navigate to `Realms > Realm Name > Services > Audit Logging`.

3. On the Secondary Configurations tab, select `Add a Secondary Configuration`. Choose JDBC from the list.

The New JDBC page appears. Enter the basic configuration for the new handler by performing the following actions:

4. Enter a name for the event handler. For example, `JDBC Audit Event Handler`.
5. In the JDBC Database URL field, enter the URL for your database server. For example, `jdbc:oracle:thin:@//host.example.com:1521/ORCL`.
6. In the JDBC Driver field, enter the classname of the driver to connect to the database. For example:
 1. `oracle.jdbc.driver.OracleDriver` - for Oracle databases
 2. `com.mysql.jdbc.Driver` - for MySQL databases
 3. `org.postgresql.Driver` - for PostgreSQL databases
7. In the Database Username field, enter the username to authenticate to the database server.

This user must have read and write privileges for the audit tables.
8. In the Database Password field, enter the password used to authenticate to the database server.
9. Enable or disable the Buffering option.
10. Select the Create button.

After the JDBC audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

11. On the General Handler Configuration tab, enable the handler and configure the topics for your audit logs:
 - a. Select Enabled to activate the event handler, if disabled.
 - b. Select the topics for your audit logs. For a description of each topic, see "Audit Log Topics".
 - c. Save your changes.
12. On the Audit Event Handler Factory tab, enter the fully qualified class name of your custom JDBC audit event handler and save your changes.
13. On the Database Configuration tab, configure the main JDBC event handler properties:
 - a. From the Database Type drop-down list, select the audit database type. The default value is `Oracle`.
 - b. In the JDBC Database URL field, enter the URL for your database server. For example, `jdbc:oracle:thin:@//host.example.com:1521/ORCL`.
 - c. In the JDBC Driver field, enter the classname of the driver to connect to the database. For example:
 1. `oracle.jdbc.driver.OracleDriver` - for Oracle databases
 2. `com.mysql.jdbc.Driver` - for MySQL databases
 3. `org.postgresql.Driver` - for PostgreSQL databases
 - d. In the Database Username field, enter the username to authenticate to the database server.
This user must have read and write privileges for the audit tables.
 - e. In the Database Password field, enter the password used to authenticate to the database server.
 - f. In the Connection Timeout (seconds) field, enter the maximum wait time before failing the connection.
 - g. In the Maximum Connection Idle Timeout (seconds) field, enter the maximum idle time in seconds before the connection is closed.
 - h. In the Maximum Connection Time (seconds) field, enter the maximum time in seconds for a connection to stay open.
 - i. In the Minimum Idle Connections field, enter the minimum number of idle connections allowed in the connection pool.
 - j. In the Maximum Connections field, enter the maximum number of connections in the connection pools.

- k. Save your changes.

14. On the Buffering tab, configure the buffering settings:

- a. Select Buffering Enabled to start audit event buffering.
- b. In the Buffer Size (number of events) field, set the size of the event buffer queue where events should queue up before being written to the database.

If the queue reaches full capacity, the process will block until a write occurs.

- c. In the Write Interval field, set the interval in seconds in which buffered events are written to the database.
- d. In the Writer Threads field, set the number of threads used to write the buffered events.
- e. In the Max Batched Events field, set the maximum number of batched statements the database can support per connection.
- f. Save your changes.

Implementing JMS Audit Event Handlers

AM supports audit logging to a JMS message broker. JMS is a Java API for sending messages between clients using a publish and subscribe model as follows:

- AM audit logging to JMS requires that the JMS message broker supports using JNDI to locate a JMS connection factory. See your JMS message broker documentation to verify that you can make connections to your broker by using JNDI before attempting to implement an AM JMS audit handler.
- AM acts as a JMS publisher client, publishing JMS messages containing audit events to a JMS *topic*.¹
- A JMS subscriber client, which is not part of the AM software and must be developed and deployed separately from AM, subscribes to the JMS topic to which AM publishes audit events. The client then receives the audit events over JMS and processes them as desired.

Before configuring the JMS audit event handler, you must perform several steps to allow AM to publish audit events as a JMS client:

¹ Note that AM and JMS use the term *topic* differently. An *AM audit topic* is a category of audit log event that has an associated one-to-one mapping to a schema type. A *JMS topic* is a distribution mechanism for publishing messages delivered to multiple subscribers.

To Prepare for JMS Audit Logging

1. Obtain JNDI connection properties that AM requires to connect to your JMS message broker. The specific connection properties vary depending on the broker. See your JMS message broker documentation for details.

For example, connecting to an Apache ActiveMQ message broker requires the following properties:

Example Apache ActiveMQ JNDI Connection Properties

Property Name	Example Value
<code>java.naming.factory.initial</code>	<code>org.apache.activemq.jndi.ActiveMQInitialContextFactory</code>
<code>java.naming.provider.url</code>	<code>tcp://localhost:61616</code>
<code>topic.audit</code>	<code>audit</code>

2. Obtain the JNDI lookup name of the JMS connection factory for your JMS message broker.
For example, for Apache ActiveMQ, the JNDI lookup name is `ConnectionFactory`.
3. Obtain the JMS client `.jar` file from your JMS message broker vendor. Add the `.jar` file to AM's classpath by placing it in the `WEB-INF/lib` directory.

For example, place the JMS client `.jar` file in `/path/to/tomcat/webapps/openam/WEB-INF/lib` if you use Apache Tomcat.

The following procedure describes how to configure a JMS audit event handler.

If your JMS message broker requires an SSL connection, you might need to perform additional, broker-dependent configuration tasks. For example, you might need to import a broker certificate into the AM keystore, or provide additional JNDI context properties.

See your JMS message broker's documentation for specific requirements for making SSL connections to your broker, and implement them as needed in addition to the steps in the following procedure.

Perform the following steps after you have installed the JMS client `.jar` file in the AM web container:

To Configure a JMS Audit Event Handler

1. Log in to the AM console as an administrator, for example `amAdmin`.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to `Configure > Global Services > Audit Logging`.

- To create the event handler in a realm, navigate to Realms > *Realm Name* > Services > Audit Logging.

3. On the Secondary Configurations tab, select Add a Secondary Configuration. Choose JMS from the list.

The New JMS configuration page appears. Enter the basic configuration for the new handler by performing the following actions:

4. Enter a name for the event handler. For example, **JMS Audit Event Handler**.
5. Select Create.

After the JMS audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

6. On the General Handler Configuration tab, enable the handler and configure the topics for your audit logs:
 - a. Select Enabled to activate the event handler, if disabled.
 - b. Select the topics for your audit logs. For a description of each topic, see "Audit Log Topics".
 - c. Save your changes.
7. On the Audit Event Handler Factory tab, keep the default class name for the audit event handler.
8. On the JMS Configuration tab, configure the main JMS event handler properties:
 - a. From the Delivery Mode drop-down list, specify the JMS delivery mode.

With persistent delivery, the JMS provider ensures that messages are not lost in transit in case of a provider failure by logging messages to storage when they are sent. Therefore, persistent delivery mode guarantees JMS message delivery, while non-persistent mode provides better performance.

The default delivery mode is **NON_PERSISTENT** delivery. Therefore, if your deployment requires delivery of every audit event to JMS subscriber clients, be sure to set the configuration to **PERSISTENT** delivery.

- b. From the Session Mode drop-down list, leave the default setting, **AUTO**, unless your JMS broker implementation requires otherwise. See your broker documentation for more information.
- c. Specify properties that AM will use to connect to your JMS message broker as key-value pairs in the JNDI Context Properties field.

AM is configured for the **audit** JNDI lookup name and JMS topic, but you can modify or delete this configuration, or add new key-value pairs. To add new key-value pairs, fill the Key and Value fields and click the Add button.

- d. In the JMS Topic Name field, specify the name of the JMS topic to which AM will publish messages containing audit events.

Subscriber clients that process AM audit events must subscribe to this topic.
 - e. In the JMS Connection Factory Name, specify the JNDI lookup name of the JMS connection factory.
 - f. Save your changes.
9. On the Batch Events tab, configure whether log events should be batched before they are published to the JMS message broker:
- a. In the Capacity field, specify the maximum capacity of the publishing queue. Execution is blocked if the queue size reaches capacity.
 - b. In the Max Batched field, specify the maximum number of events to be delivered when AM publishes the events to the JMS message broker.
 - c. In the Writing Interval field, specify the interval (in seconds) between transmissions to JMS.
 - d. Save your changes.

Implementing Splunk Audit Event Handlers

AM supports sending audit logging data to a Splunk HTTP event collector REST endpoint. This endpoint requires an authentication token created in Splunk and configured in the AM event handler.

Before configuring the Splunk audit event handler in AM, configure the endpoint in Splunk:

To Prepare Splunk

Perform the following steps in Splunk:

1. Create a source type to match AM audit logs. Consider the following configuration tips:
 - It must be a structured type.
 - It must not have indexed extractions.
 - It must have event breaks on every line.
 - Timestamp extraction must be automatic.
2. Create an HTTP event collector token that uses the source type you just created.
3. Obtain the HTTP event collector token value for the token you just created. For example, **AD565CB5-BB8A-40FD-8A7C-94F549CEDEC**F.

This token, which allows AM to log data to Splunk, is required for the AM audit event handler configuration.

4. Obtain the HTTP event collector port. For example, **8088**.
5. Ensure that the HTTP event collector and its tokens are enabled.

The following procedure describes how to configure a Splunk audit event handler in AM. Perform the following steps after you have created a Splunk HTTP event collector token:

To Configure a Splunk Audit Event Handler

1. Log in to the AM console as an administrator, for example **amAdmin**.
2. Determine whether to create the event handler in a realm or use the default global event handler, then take one of the following actions:
 - To create the event handler in the global configuration, navigate to **Configure > Global Services > Audit Logging**.
 - To create the event handler in a realm, navigate to **Realms > *Realm Name* > Services > Audit Logging**.

3. On the Secondary Configurations tab, select **Add a Secondary Configuration**. Choose Splunk from the list.

The New Splunk configuration page appears. Enter the basic configuration for the new handler by performing the following actions:

4. Enter a name for the event handler. For example, **Splunk Audit Event Handler**.
5. In the Server Hostname field, enter the hostname or IP address of the Splunk HTTP event collector to which AM should connect when writing audit events.
6. In the Server Port field, enter the port number to access the Splunk HTTP event collector.
7. In the Authorization Token field, enter the authorization token for the Splunk HTTP event collector REST endpoint.
8. Select **Create**.

After the Splunk audit event handler is created, several configuration tabs appear. To configure the event handler, perform the following steps:

9. On the General Handler Configuration tab, enable the handler and configure the topics for your audit logs:
 - a. Select **Enabled** to activate the event handler, if disabled.

- b. Select the topics for your audit logs. For a description of each topic, see "Audit Log Topics".
 - c. Save your changes.
10. On the Audit Event Handler Factory tab, keep the default class name for the audit event handler.
11. On the Splunk Configuration tab, configure the main Splunk event handler properties:
 - a. In the Server Hostname field, enter the hostname or IP address of the Splunk server to which AM should connect when writing audit events.
 - b. In the Server Port field, enter the port number to access the Splunk HTTP event collector REST endpoint.
 - c. If SSL is enabled and configured between AM and the Splunk server, select SSL Enabled.
 - d. In the Authorization Token field, enter the authorization token for the Splunk HTTP event collector REST endpoint.
 - e. Save your changes.
12. On the Buffering tab, configure options about how log events should be buffered in memory before they are written to the Splunk data store:
 - a. In the Batch Size field, specify the number of audit events that AM pulls from the audit buffer when writing a batch of events to Splunk.

When buffering is enabled, all audit events are put into an in-memory buffer (one per handled topic), so that the original thread that generated the event can fulfill the requested operation, rather than wait for I/O to complete. A dedicated thread (one per handled topic) constantly pulls events from the buffer in batches and writes them to Splunk. If the buffer becomes empty, the dedicated thread goes to sleep until a new item gets added.
 - b. In the Queue Capacity field, specify the maximum number of audit events that AM can queue in this audit handler's buffer.

If the number of events to queue exceeds the queue capacity, AM raises an exception and the excess audit events are dropped, and therefore not written to Splunk.
 - c. In the Write interval (in milliseconds) field, specify how often AM should write buffered events to Splunk.
 - d. Save your changes.

Configuring the Trust Transaction Header System Property

AM supports the propagation of the transaction ID across the ForgeRock platform, such as from DS or IDM to AM, using the HTTP header `X-ForgeRock-TransactionId`. The `X-ForgeRock-TransactionId` header

is automatically set in all outgoing HTTP calls from one ForgeRock product to another. Customers can also set this header themselves from their own applications or scripts calling into the ForgeRock platform.

You can set a new property `org.forgerock.http.TrustTransactionHeader` to `true`, which will trust any incoming `X-ForgeRock-TransactionId` headers. By default, the `org.forgerock.http.TrustTransactionHeader` is set to `false`, so that a malicious actor cannot flood the system with requests using the same transaction ID header to hide their tracks.

To Configure the Trust Transactions Header System Property

1. Log in to the AM console.
2. Navigate to Configure > Server Defaults > Advanced.
3. In the Add a Name field, enter `org.forgerock.http.TrustTransactionHeader`, and enter `true` in the corresponding Add a Value field.
4. Save your changes.

Your AM instance will now accept incoming `X-ForgeRock-TransactionId` headers, which can then be tracked in the audit logs.

5. Repeat this procedure for all servers requiring this property.

Implementing the Classic Logging Service

Note

AM supports two Audit Logging Services: the classic Logging Service, which is based on a Java SDK and is available in AM versions prior to AM 5.0, and a common REST-based Audit Logging Service. The classic Logging Service is deprecated.

To configure AM logging properties, log in to the AM console as AM administrator, and navigate to Configure > Global Services > Logging.

For more information on the available settings, see "Audit Logging" in the *Reference* reference.

Audit Logging to Flat Files

By default, AM audit logs are written to files in the configuration directory for the instance, such as `$HOME/openam/log/`.

AM sends messages to different log files, each named after the service logging the message, with two different types log files per service: `.access` and `.error`. Thus, the current log files for the authentication service are named `amAuthentication.access` and `amAuthentication.error`.

For details, see "Log Files and Messages" in the *Reference*.

Audit Logging to a Syslog Server

AM supports sending audit log messages to a syslog server for collation.

You can enable syslog audit logging by using the AM console, or the `ssoadm` command.

Enabling Syslog Audit Logging by Using the AM Console

1. Log in to the AM console as AM administrator, for example `amAdmin`.
2. Navigate to Configure > Global Services > Logging.
3. On the Syslog tab, configure the following settings as appropriate for your syslog server, and save your changes:

- Syslog server host
- Syslog server port
- Syslog server protocol
- Syslog facility
- Syslog connection timeout

For information on these settings, see "Logging" in the *Reference*.

4. On the General tab, set the Logging Type drop-down list to `Syslog`, and save your changes.

Enabling Syslog Audit Logging by Using the `ssoadm` Command

1. Create a text file, for example, `MySyslogServerSettings.txt` containing the settings used when audit logging to a syslog server, as shown below:

```
iplanet-am-logging-syslog-port=514
iplanet-am-logging-syslog-protocol=UDP
iplanet-am-logging-type=Syslog
iplanet-am-logging-syslog-connection-timeout=30
iplanet-am-logging-syslog-host=localhost
iplanet-am-logging-syslog-facility=local5
```

2. Use the following SSOADM command to configure audit logging to a syslog server:

```
$ ssoadm \
  set-attr-defs \
    --adminid uid=amAdmin,ou=People,dc=openam,dc=forgerock,dc=org \
    --password-file /tmp/pwd.txt \
    --servicename iPlanetAMLoggingService \
    --schematype Global \
    --datafile MySyslogServerSettings.txt
Schema attribute defaults were set.
```

Chapter 12

Reference

This reference section covers other information relating to securing an AM instance. For the global services reference, see [Reference](#).

Reference Entries About...	
Audit Logging Learn more about the log format of the different files and tables used by the audit logging service.	See "Audit Logging Reference".
Session Learn how to customize CTS-based session quota exhaust actions.	See "Customizing CTS-Based Session Quota Exhaust Actions".

Audit Logging Reference

AM writes log messages generated from audit events triggered by its components, instances, and other ForgeRock-based stack products.

Audit Log Format

This section presents the audit log format for each topic-based file, event names, and audit constants used in its log messages.

Access Log Format

Access Log Format

Schema Property	Description
<code>_id</code>	Specifies a universally unique identifier (UUID) for the message object, such as <code>a568d4fe-d655-49a8-8290-bfc02095bec9-491</code> .
<code>timestamp</code>	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: <code>yyyy-MM-ddTHH:mm:ss.msZ</code> . For example: <code>2015-11-14T00:16:04.653Z</code>
<code>eventName</code>	Specifies the name of the audit event. For example, <code>AM-ACCESS-ATTEMPT</code> and <code>AM-ACCESS-OUTCOME</code> . For a list of audit event names, see "Audit Log Event Names".

Schema Property	Description
<code>transactionId</code>	<p>Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID even for different audit event topics. For example, <code>9c9e8d5c-2941-4e61-9c3c-8a990088e801</code>.</p> <p>AM supports a feature where trusted AM deployment with multiple instances, components, and ForgeRock stack products can propagate the transaction ID through each call across the stack. AM reads the <code>X-ForgeRock-TransactionId</code> HTTP header and appends an integer to the transaction ID. Note that this feature is disabled by default. When enabled, this feature should filter the <code>X-ForgeRock-TransactionId</code> HTTP header for connections from untrusted sources.</p>
<code>user.id</code>	Specifies the universal identifier for authenticated users. For example, <code>id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com</code> .
<code>trackingIds</code>	<p>Specifies a unique random string generated as an alias for each AM session ID and OAuth 2.0 token. In releases prior to OpenAM 13.0.0, the <code>contextId</code> log property used a random string as an alias for the session ID. The <code>trackingIds</code> property also uses an alias when referring to session IDs, for example, <code>["45b17894529cf74301"]</code>.</p> <p>OpenAM 13.0.0 extended this property to handle OAuth 2.0 tokens. In this case, whenever AM generates an access or grant token, it also generates unique random value and logs it as an alias. In this way, it is possible to trace back an access token back to its originating grant token, trace the grant token back to the session in which it was created, and then trace how the session was authenticated. An example of a <code>trackingIds</code> property in an OAuth 2.0/ OpenID Connect 1.0 environment is: <code>["1979edf68543ead001", "8878e51a-f2aa-464f-b1cc-b12fd6daa415", "3df9a5c3-8d1e-4ee3-93d6-b9bbe58163bc"]</code></p>
<code>server.ip</code>	Specifies the IP address of the AM server. For example, <code>127.0.0.1</code> .
<code>server.port</code>	Specifies the port number used by the AM server. For example, <code>8080</code> .
<code>client.host</code>	Specifies the client hostname. This field is only populated if reverse DNS lookup is enabled.
<code>client.ip</code>	Specifies the client IP address.
<code>client.port</code>	Specifies the client port number.
<code>authorizationId.roles</code>	Specifies the list of roles for the authorized user.
<code>authorizationId.component</code>	Specifies the component part of the authorized ID, such as
<code>request.protocol</code>	Specifies the protocol associated with the request operation. Possible values: <code>CREST</code> and <code>PLL</code> .
<code>request.operation</code>	<p>Specifies the request operation. For Common REST operations, possible values are: <code>READ</code>, <code>ACTION</code>, <code>QUERY</code>.</p> <p>For PLL operations, possible values are: <code>LoginIndex</code>, <code>SubmitRequirements</code>, <code>GetSession</code>, <code>REQUEST_ADD_POLICY_LISTENER</code>.</p>
<code>request.detail</code>	Specifies the detailed information about the request operation. For example:

Schema Property	Description
	<ul style="list-style-type: none"> • {"action": "idFromSession"} • {"action": "validateGoto"} • {"action": "validate"} • {"action": "logout"} • {"action": "schema"} • {"action": "template"}
http.method	Specifies the HTTP method requested by the client. For example, GET, POST, PUT.
http.path	Specifies the path of the HTTP request. For example, https://openam.example.com:8443/openam/json/realms/root/authenticate.
http.queryParameters	Specifies the HTTP query parameter string. For example: <ul style="list-style-type: none"> • { "_action": ["idFromSession"] } • { "_queryFilter": ["true"] } • { "_action": ["validate"] } • { "_action": ["logout"] } • { "realm": ["/shop"] } • { "_action": ["validateGoto"] }
http.request.headers	Specifies the HTTP header for the request. For example: <pre>{ "accept":["application/json, text/javascript, */*; q=0.01"], "Accept-API-Version":["protocol=1.0"], "accept-encoding":["gzip, deflate"], "accept-language":["en-US;q=1,en;q=0.9"], "cache-control":["no-cache"], "connection":["Keep-Alive"], "content-length":["0"], "host":["forgerock-am.openrock.org"] }</pre>

Schema Property	Description
	<pre>], "pragma": ["no-cache"], "referer": ["https://forgerock-am.openrock.org/openam/XUI/"], "user-agent": ["Mozilla/5.0 (X11; Linux x86_64; rv:31.0) Gecko/20100101 Firefox/31.0"], "x-nosession": ["true"], "x-requested-with": ["XMLHttpRequest"], "x-username": ["anonymous"] } </pre> <p>Note: line feeds and truncated values in the example are for readability purposes.</p>
<code>http.request.cookies</code>	Specifies a JSON map of key-value pairs and appears as its own property to allow for blacklisting fields or values.
<code>http.response.cookies</code>	Not used in AM.
<code>response.status</code>	Specifies the response status of the request. For example, <code>SUCCESS</code> , <code>FAILURE</code> , or <code>null</code> .
<code>response.statusCode</code>	Specifies the response status code, depending on the protocol. For Common REST, HTTP failure codes are displayed but not HTTP success codes. For PLL endpoints, PLL error codes are displayed.
<code>response.detail</code>	Specifies the message associated with <code>response.statusCode</code> . For example, the <code>response.statusCode</code> of <code>401</code> has a <code>response.detail</code> of <code>{ "reason": "Unauthorized" }</code> .
<code>response.elapsedTime</code>	Specifies the time to execute the access event, usually in millisecond precision.
<code>response.elapsedTimeUnits</code>	Specifies the elapsed time units of the response. For example, <code>MILLISECONDS</code> .
<code>component</code>	Specifies the AM service utilized. For example, <code>Server Info</code> , <code>Users</code> , <code>Config</code> , <code>Session</code> , <code>Authentication</code> , <code>Policy</code> , <code>OAuth</code> , <code>Web Policy Agent</code> , or <code>Java Policy Agent</code> .
<code>realm</code>	Specifies the realm where the operation occurred. For example, the Top Level Realm (<code>"/"</code>) or the sub-realm name (<code>"/shop"</code>).

Activity Log Format

Activity Log Format

Property	Description
<code>_id</code>	Specifies a universally unique identifier (UUID) for the message object, such as <code>a568d4fe-d655-49a8-8290-bfc02095bec9-487</code> .
<code>timestamp</code>	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: <code>yyyy-MM-ddTHH:mm:ss.msZ</code> . For example: <code>2015-11-14T00:16:04.652Z</code>
<code>eventName</code>	Specifies the name of the audit event. For example, <code>AM-SESSION_CREATED</code> , <code>AM-SESSION-LOGGED_OUT</code> , <code>AM-IDENTITY-CHANGE</code> . For a list of audit event names, see "Audit Log Event Names".
<code>transactionId</code>	Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for same even for different audit event topics. For example, <code>9c9e8d5c-2941-4e61-9c3c-8a990088e801</code> .
<code>user.id</code>	Specifies the universal identifier for authenticated users. For example, <code>id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com</code> .
<code>trackingIds</code>	Specifies an array containing a random context ID that identifies the session and a random string generated from an OAuth 2.0/OpenID Connect 1.0 flow that could track an access token ID or an grant token ID. For example, <code>["45b17894529cf74301"]</code> .
<code>runAs</code>	Specifies the user to run the activity as. May be used in delegated administration. For example, <code>id=dsameuser,ou=user,dc=example,dc=com</code> .
<code>objectId</code>	Specifies the identifier of an object that has been created, updated, or deleted. For logging sessions, the session <code>trackingId</code> is used in this field. For example, <code>["45b17894529cf74301"]</code>
<code>operation</code>	Specifies the state change operation invoked: <code>CREATE</code> , <code>MODIFY</code> , or <code>DELETE</code> .
<code>before</code>	Not used.
<code>after</code>	Not used.
<code>changedFields</code>	Not used.
<code>revision</code>	Not used.
<code>component</code>	Specifies the AM service utilized. For example, <code>Session</code> or <code>Self-Service</code> .
<code>realm</code>	Specifies the realm where the operation occurred. For example, the Top Level Realm (<code>"/</code>) or the sub-realm name (<code>"/shop</code>).

Authentication Log Format

Authentication Log Format

Property	Description
<code>_id</code>	Specifies a universally unique identifier (UUID) for the message object, such as <code>a568d4fe-d655-49a8-8290-bfc02095bec9-485</code> .
<code>timestamp</code>	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: <code>yyyy-MM-ddTHH:mm:ss.msZ</code> . For example: <code>2015-11-14T00:16:04.640Z</code>
<code>eventName</code>	Specifies the name of the audit event. For example, <code>AM-LOGOUT</code> and <code>AM-LOGIN-MODULE-COMPLETED</code> . For a list of audit event names, see "Audit Log Event Names".
<code>transactionId</code>	Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for same even for different audit event topics. For example, <code>9c9e8d5c-2941-4e61-9c3c-8a990088e801</code> .
<code>user.id</code>	Specifies the universal identifier for authenticated users. For example, <code>id=scarter, ou=user, o=shop, ou=services, dc=example, dc=com</code> .
<code>trackingIds</code>	Specifies an array containing a unique random context ID. For example: <ul style="list-style-type: none"> For OAuth 2.0/OpenID Connect flows, it identifies the session and a random string generated that can track an access token ID or a grant token ID. For authentication trees, it identifies an authentication tree flow.
<code>result</code>	Depending on the event being logged, specifies the outcome of: <ul style="list-style-type: none"> A single authentication module within a chain The result for an authentication tree Possible values are <code>SUCCESSFUL</code> or <code>FAILED</code> .
<code>principal</code>	Specifies the array of accounts used to authenticate, such as <code>["amadmin"]</code> and <code>["scarter"]</code> .
<code>context</code>	Not used
<code>entries</code>	Specifies the JSON representation of the details of an authentication module, chain, tree or node. AM creates an event as each module or node completes and a final event at the end of the chain or tree. Examples:

Property	Description
	<pre> "entries":[{ "moduleId":"DataStore", "info":{ "moduleClass":"DataStore", "ipAddress":"127.0.0.1", "moduleName":"DataStore", "authLevel":"0" } }] </pre>
	<pre> "entries":[{ "info":{ "nodeOutcome":"true", "treeName":"Example", "displayName":"Data Store Decision", "nodeType":"DataStoreDecisionNode", "nodeId":"e5ec495a-2ae2-4eca-8afb-9781dea04170", "authLevel":"0" } }] </pre>
component	Specifies the AM service utilized. For example, Authentication .
realm	Specifies the realm where the operation occurred. For example, the Top Level Realm ("/") or the sub-realm name ("/shop").

Config Log Format

Config Log Format

Property	Description
_id	Specifies a universally unique identifier (UUID) for the message object. For example, 6a568d4fe-d655-49a8-8290-bfc02095bec9-843 .
timestamp	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: yyyy-MM-ddTHH:mm:ss.msZ . For example, 2015-11-14T00:21:03.490Z .
eventName	Specifies the name of the audit event. For example, AM-CONFIG-CHANGE . For a list of audit event names, see "Audit Log Event Names".
transactionId	Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for different audit event topics. For example, 301d1a6e-67f9-4e45-bfeb-5e4047a8b432 .

Property	Description
<code>user.id</code>	Not used. You can determine the value for this field by linking to the access event using the same <code>transactionId</code> .
<code>trackingIds</code>	Not used.
<code>runAs</code>	Specifies the user to run the activity as. May be used in delegated administration. For example, <code>uid=amAdmin,ou=People,dc=example,dc=com</code> .
<code>objectId</code>	Specifies the identifier of a system object that has been created, modified, or deleted. For example, <code>ou=SamuelTwo,ou=default,ou=OrganizationConfig,ou=1.0,ou=iPlanetAMAuthSAML2Service,ou=services,o=shop,ou=services,dc=example,dc=com</code> .
<code>operation</code>	Specifies the state change operation invoked: <code>CREATE</code> , <code>MODIFY</code> , or <code>DELETE</code> .
<code>before</code>	Specifies the JSON representation of the object prior to the activity. For example: <pre> { "sunsmpriority":["0"], "objectclass":["top", "sunServiceComponent", "organizationalUnit"], "ou":["SamuelTwo"], "sunserviceID":["serverconfig"] } </pre>
<code>after</code>	Specifies the JSON representation of the object after the activity. For example:

Property	Description
	<pre>{ "sunKeyValue": ["forgerock-am-auth-saml2-auth-level=0", "forgerock-am-auth-saml2-meta-alias=/sp", "forgerock-am-auth-saml2-entity-name=http://", "forgerock-am-auth-saml2-authn-context-decl-ref=", "forgerock-am-auth-saml2-force-authn=none", "forgerock-am-auth-saml2-is-passive=none", "forgerock-am-auth-saml2-login-chain=", "forgerock-am-auth-saml2-auth-comparison=none", "forgerock-am-auth-saml2-req-binding=urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect", "forgerock-am-auth-saml2-binding=urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact", "forgerock-am-auth-saml2-authn-context-class-ref=", "forgerock-am-auth-saml2-slo-relay=http://", "forgerock-am-auth-saml2-allow-create=false", "forgerock-am-auth-saml2-name-id-format=urn:oasis:names:tc:SAML:2.0:nameid-format:persistent"] }</pre>
changedFields	Specifies the fields that were changed. For example, ["sunKeyValue"].
revision	Not used.
component	Not used.
realm	Specifies the realm where the operation occurred. For example, the Top Level Realm ("/") or the sub-realm name ("/shop").

Audit Log Event Names

The following section presents the predefined names for the audit events:

Audit Log Event Names

Topic	EventName
access	AM-ACCESS_ATTEMPT
access	AM-ACCESS-OUTCOME
activity	AM-SELFERVICE-REGISTRATION-COMPLETED
activity	AM-SELFERVICE-PASSWORDCHANGE-COMPLETED
activity	AM-SESSION-CREATED
activity	AM-SESSION-IDLE_TIME_OUT
activity	AM-SESSION-MAX_TIMED_OUT
activity	AM-SESSION-LOGGED_OUT

Topic	EventName
activity	AM-SESSION-DESTROYED
activity	AM-SESSION-PROPERTY_CHANGED
activity	AM-IDENTITY-CHANGE
activity	AM-GROUP-CHANGE
authentication	AM-LOGOUT
authentication	AM-LOGIN-COMPLETED
authentication	AM-LOGIN-MODULE-COMPLETED
authentication	AM-NODE-LOGIN-COMPLETED
authentication	AM-TREE-LOGIN-COMPLETED
config	AM-CONFIG-CHANGE

Audit Log Components

The following section presents the predefined audit event components that make up the log messages:

Audit Log Event Components

Event Component	AM Component, Service, or Feature
OAuth	OAuth 2.0, OpenID Connect 1.0, and UMA
CTS	Core Token Service
AM Agents	Web and Java agents
Authentication	Authentication service
Dashboard	Dashboard service
Server Info	Server information service
Users	Users component
Groups	Groups component
Oath	Mobile authentication
Devices	Trusted devices
Policy	Policies
Realms	Realms and sub-realms
Session	Session service
Script	Scripting service
Batch	Batch service
Config	Configuration

Event Component	AM Component, Service, or Feature
STS	Secure Token Service: REST and SOAP
Record	Recording service
Audit	Auditing service
Radius	RADIUS server
Self-Service	User Self-Service service
ssoadm	ssoadm command
SAML2	SAML v2.0
Push	Push Notification service

Audit Log Failure Reasons

The following section presents the predefined audit event failure reasons:

Audit Log Event Authentication Failure Reasons

Failure	Description
LOGIN_FAILED	Incorrect/invalid credentials presented.
INVALID_PASSWORD	Invalid credentials entered.
NO_CONFIG	Authentication chain does not exist.
NO_USER_PROFILE	No user profile found for this user.
USER_INACTIVE	User is not active.
LOCKED_OUT	Maximum number of failure attempts exceeded. User is locked out.
ACCOUNT_EXPIRED	User account has expired.
LOGIN_TIMEOUT	Login timed out.
MODULE_DENIED	Authentication module is denied.
MAX_SESSION_REACHED	Limit for maximum number of allowed sessions has been reached.
INVALID_REALM	Realm does not exist.
REALM_INACTIVE	Realm is not active.
USER_NOTE_FOUND	Role-based authentication: user does not belong to this role.
AUTH_TYPE_DENIED	Authentication type is denied.
SESSION_CREATE_ERROR	Cannot create a session.
INVALID_LEVEL	Level-based authentication: Invalid authentication level.

Audit Log Default Whitelist

When an object is passed in an audit event, it might contain information that should not be logged. By default, the AM uses a whitelist to specify which fields of the event appear.

The following fields appear on the default, built in whitelist. This lists specifies each field by its JSON path. If a whitelisted field contains an object, then listing the field means the whole object is whitelisted:

Default *Access* Log Whitelist

- `/_id`
- `/client`
- `/eventName`
- `/http/request/headers/accept`
- `/http/request/headers/accept-api-version`
- `/http/request/headers/content-type`
- `/http/request/headers/host`
- `/http/request/headers/user-agent`
- `/http/request/headers/x-forwarded-for`
- `/http/request/headers/x-forwarded-host`
- `/http/request/headers/x-forwarded-port`
- `/http/request/headers/x-forwarded-proto`
- `/http/request/headers/x-original-uri`
- `/http/request/headers/x-real-ip`
- `/http/request/headers/x-request-id`
- `/http/request/headers/x-requested-with`
- `/http/request/headers/x-scheme`
- `/http/request/method`
- `/http/request/path`
- `/http/request/queryParameters/authIndexType`
- `/http/request/queryParameters/authIndexValue`
- `/http/request/queryParameters/composite_advice`

- /http/request/queryParameters/level
- /http/request/queryParameters/module_instance
- /http/request/queryParameters/resource
- /http/request/queryParameters/role
- /http/request/queryParameters/service
- /http/request/queryParameters/user
- /http/request/secure
- /request
- /response
- /server
- /timestamp
- /trackingIds
- /transactionId
- /userId

Default **Activity** Log Whitelist

- /_id
- /after/assignedDashboard
- /after/cn
- /after/commonName
- /after/givenName
- /after/inetUserStatus
- /after/iplanet-am-user-alias-list
- /after/iplanet-am-user-login-status
- /after/kbaInfoAttempts
- /after/memberof
- /after/o
- /after/oath2faEnabled
- /after/objectClass

- /after/organizationName
- /after/organizationUnitName
- /after/ou
- /after/push2faEnabled
- /after/sn
- /after/sunAMAuthInvalidAttemptsData
- /after/surname
- /after/uid
- /after/uniqueMember
- /after/userid
- /before/assignedDashboard
- /before/cn
- /before/commonName
- /before/givenName
- /before/inetUserStatus
- /before/iplanet-am-user-alias-list
- /before/iplanet-am-user-login-status
- /before/kbaInfoAttempts
- /before/memberof
- /before/o
- /before/oath2faEnabled
- /before/objectClass
- /before/organizationName
- /before/organizationUnitName
- /before/ou
- /before/push2faEnabled
- /before/sn
- /before/sunAMAuthInvalidAttemptsData

- /before/surname
- /before/uid
- /before/uniqueMember
- /before/userid
- /changedFields
- /component
- /component
- /eventName
- /objectId
- /operation
- /realm
- /realm
- /revision
- /runAs
- /timestamp
- /trackingIds
- /transactionId
- /userId

Default *Authentication* Log Whitelist

- /

Default *Config* Log Whitelist

- /_id
- /changedFields
- /component
- /eventName
- /objectId
- /operation

- /realm
- /revision
- /runAs
- /timestamp
- /trackingIds
- /transactionId
- /userId

JDBC Audit Log Tables

AM writes audit events to relational databases using the JDBC audit event handler. This section presents the columns for each audit table.

am_auditaccess

am_auditaccess

Column	Datatype	Description
id	VARCHAR(56) NOT NULL	Specifies a universally unique identifier (UUID) for the message object, such as <code>a568d4fe-d655-49a8-8290-bfc02095bec9-491</code> .
timestamp_	VARCHAR(29) NULL	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: <code>yyyy-MM-ddTHH:mm:ss.msZ</code> . For example: <code>2015-11-14T00:16:04.653Z</code>
transactionid	VARCHAR(255) NULL	<p>Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for different audit event topics. For example, <code>9c9e8d5c-2941-4e61-9c3c-8a990088e801</code>.</p> <p>AM supports a feature where a trusted AM deployment with multiple instances, components, and ForgeRock products can propagate a transaction ID through each call across the stack. AM reads the <code>X-ForgeRock-TransactionId</code> HTTP header and appends an integer to the transaction ID. Note that this feature is disabled by default. When enabled, this feature should filter the <code>X-ForgeRock-TransactionId</code> HTTP header for connections from untrusted sources.</p>
eventname	VARCHAR(255)	Specifies the name of the audit event. For example, <code>AM-ACCESS-ATTEMPT</code> and <code>AM-ACCESS-OUTCOME</code> . For a list of audit event names, see "Audit Log Event Names".

Column	Datatype	Description
userid	VARCHAR(255) NULL	Specifies the universal identifier for the authenticated user. For example, <code>id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com</code> .
trackingids	MEDIUMTEXT	Specifies the tracking IDs of the event, used by all topics.
server_ip	VARCHAR(40)	Specifies the IP address of the AM server.
server_port	VARCHAR(5)	Specifies the port number used by the AM server. For example, <code>8080</code> .
client_host	VARCHAR(255)	Specifies the client hostname. This column is only populated if reverse DNS lookup is enabled.
client_ip	VARCHAR(40)	Specifies the client IP address.
client_port	VARCHAR(5)	Specifies the client port number.
request_protocol	VARCHAR(255) NULL	Specifies the protocol associated with the request operation. Possible values: <code>CREST</code> and <code>PLL</code> .
request_operation	VARCHAR(255) NULL	Specifies the request operation. For Common REST operations, possible values: <code>READ</code> , <code>ACTION</code> , <code>QUERY</code> . For PLL operations, possible values: <code>LoginIndex</code> , <code>SubmitRequirements</code> , <code>GetSession</code> , <code>REQUEST_ADD_POLICY_LISTENER</code> .
request_detail	TEXT NULL	Specifies the detailed information about the request operation. For example: <ul style="list-style-type: none"> <code>{"action": "idFromSession"}</code> <code>{"action": "validateGoto"}</code> <code>{"action": "validate"}</code> <code>{"action": "logout"}</code> <code>{"action": "schema"}</code> <code>{"action": "template"}</code>
http_request_secure	BOOLEAN NULL	Specifies the HTTP method requested by the client. For example, <code>true</code> or <code>false</code> . Note that <code>false</code> does not mean the client connection is insecure as there may be a reverse proxy terminating the HTTPS connection.
http_request_method	VARCHAR(7) NULL	Specifies the HTTP method requested by the client. For example, <code>GET</code> , <code>POST</code> , <code>PUT</code> .
http_request_path	VARCHAR(255) NULL	Specifies the path of the HTTP request. For example, <code>https://openam.example.com:8443/openam/json/realms/root/authenticate</code> .
http_request_queryparameters	MEDIUMTEXT NULL	Specifies the HTTP query parameter string. For example: <ul style="list-style-type: none"> <code>{ "_action": ["idFromSession"] }</code> <code>{ "_queryFilter": ["true"] }</code>

Column	Datatype	Description
		<ul style="list-style-type: none"> • { "_action": ["validate"] } • { "_action": ["logout"] } • { "realm": ["/shop"] } • { "_action": ["validateGoto"] }
http_request_headers	MEDIUMTEXT NULL	<p>Specifies the HTTP headers for the request. For example:</p> <pre> { "accept":["application/json, text/javascript, */*; q=0.01"], "Accept-API-Version":["protocol=1.0"], "accept-encoding":["gzip, deflate"], "accept-language":["en-US;q=1,en;q=0.9"], "cache-control":["no-cache"], "connection":["Keep-Alive"], "content-length":["0"], "host":["forgerock-am.openrock.org"], "pragma":["no-cache"], "referer":["https://forgerock-am.openrock.org/openam/XUI/"], "user-agent":["Mozilla/5.0 (X11; Linux x86_64; rv:31.0) Gecko/20100101 Firefox/31.0"], "x-nosession":["true"], "x-requested-with":["XMLHttpRequest"], "x-username":["anonymous"] }</pre>

Column	Datatype	Description
		Note: line feeds and truncated values in the example are for readability purposes.
http_request_cookies	MEDIUMTEXT NULL	Specifies a JSON map of key-value pairs and appears as its own property to allow for blacklisting fields or values. For example: <pre>"cookies": "amlbcookie=01; iPlanetDirectoryPro=\"AQIC5wM2LY...*AAJTSQACMfwT...*\"; iPlanetDirectoryPro=eyJ0eXAiOiJK...eyJzdWIiOiJkZ..."</pre> Note: line feeds and truncated values in the example are for readability purposes.
http_response_headers	MEDIUMTEXT NULL	Captures the headers returned by AM to the client (that is, the inverse of <code>http_request_headers</code> . Note that AM does not currently populate this field.
response_status	VARCHAR(10) NULL	Specifies the response status of the request. For example, <code>SUCCESS</code> , <code>FAILURE</code> , <code>ALLOWED</code> , <code>DENIED</code> , or <code>NULL</code> .
response_statuscode	VARCHAR(255) NULL	Specifies the response status code, depending on the protocol. For Common REST, HTTP failure codes are displayed but not HTTP success codes. For PLL endpoints, PLL error codes are displayed.
response_detail	TEXT NULL	Specifies the message associated with the response status code. For example, a response status code of 401 has a response detail of <code>{ "reason": "Unauthorized" }</code> .
response_elapsedtime	VARCHAR(255) NULL	Specifies the time to execute the access event, usually in millisecond precision.
response_elapsedtimeunits	VARCHAR(255) NULL	Specifies the elapsed time units of the response. For example, <code>MILLISECONDS</code> .
component	VARCHAR(255) NULL	Specifies the AM service utilized. For example, <code>Server Info</code> , <code>Users</code> , <code>Config</code> , <code>Session</code> , <code>Authentication</code> , <code>Policy</code> , <code>OAuth</code> .
realm	VARCHAR(255) NULL	Specifies the realm where the operation occurred. For example, the Top Level Realm <code>("/")</code> or the sub-realm name <code>("/shop")</code> .

am_auditauthentication

am_auditauthentication

Column	Datatype	Description
id	VARCHAR(56) NOT NULL	Specifies a universally unique identifier (UUID) for the message object, such as <code>a568d4fe-d655-49a8-8290-bfc02095bec9-491</code> .
timestamp_	VARCHAR(29) NULL	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: <code>yyyy-MM-ddTHH:mm:ss.msZ</code> . For example: <code>2015-11-14T00:16:04.653Z</code>
transactionid	VARCHAR(255) NULL	Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any

Column	Datatype	Description
		<p>events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for different audit event topics. For example, <code>9c9e8d5c-2941-4e61-9c3c-8a990088e801</code>.</p> <p>AM supports a feature where a trusted AM deployment with multiple instances, components, and ForgeRock products can propagate a transaction ID through each call across the stack. AM reads the <code>X-ForgeRock-TransactionId</code> HTTP header and appends an integer to the transaction ID. Note that this feature is disabled by default. When enabled, this feature should filter the <code>X-ForgeRock-TransactionId</code> HTTP header for connections from untrusted sources.</p>
eventname	VARCHAR(255) NULL	Specifies the name of the audit event. For example, <code>AM-LOGIN-MODULE-COMPLETED</code> and <code>AM-LOGOUT</code> . For a list of audit event names, see "Audit Log Event Names".
userid	VARCHAR(255) NULL	Specifies the universal identifier for authenticated users. For example, <code>id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com</code> .
trackingids	MEDIUMTEXT	Specifies the tracking IDs of the event, used by all topics.
result	VARCHAR(255) NULL	<p>Depending on the event being logged, specifies the outcome of:</p> <ul style="list-style-type: none"> • A single authentication module within a chain • The result for an authentication tree <p>Possible values are <code>SUCCESSFUL</code> or <code>FAILED</code>.</p>
principals	MEDIUMTEXT	Specifies the array of accounts used to authenticate, such as <code>["amadmin"]</code> and <code>["scarter"]</code> .
context	N/A	MEDIUMTEXT. Not used.
entries	MEDIUMTEXT	<p>Specifies the JSON representation of the details of an authentication module, chain, tree or node. AM creates an event as each module or node completes and a final event at the end of the chain or tree. For example:</p> <pre> "entries":[{ "moduleId":"DataStore", "info":{ "moduleClass":"DataStore", "ipAddress":"127.0.0.1", "moduleName":"DataStore", "authLevel":"0" } }] </pre>

Column	Datatype	Description
		<pre> "entries":[{ "info":{ "nodeOutcome":"true", "treeName":"Example", "displayName":"Data Store Decision", "nodeType":"DataStoreDecisionNode", "nodeId":"e5ec495a-2ae2-4eca-8afb-9781dea04170", "authLevel":"0" } }] </pre>
component	VARCHAR(255) NULL	Specifies the AM service utilized. For example, Server Info, Users, Config, Session, Authentication, Policy, OAuth .
realm	VARCHAR(255) NULL	Specifies the realm where the operation occurred. For example, the Top Level Realm ("/") or the sub-realm name ("/shop").

am_auditactivity

am_auditactivity

Column	Datatype	Description
id	VARCHAR(56) NOT NULL	Specifies a universally unique identifier (UUID) for the message object, such as a568d4fe-d655-49a8-8290-bfc02095bec9-491 .
timestamp_	VARCHAR(29) NOT NULL	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: yyyy-MM-ddTHH:mm:ss.msZ . For example: 2015-11-14T00:16:04.653Z
transactionid	VARCHAR(255) NULL	<p>Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for different audit event topics. For example, 9c9e8d5c-2941-4e61-9c3c-8a990088e801.</p> <p>AM supports a feature where a trusted AM deployment with multiple instances, components, and ForgeRock products can propagate a transaction ID through each call across the stack. AM reads the X-ForgeRock-TransactionId HTTP header and appends an integer to the transaction ID. Note that this feature is disabled by default. When enabled, this feature should filter the X-ForgeRock-TransactionId HTTP header for connections from untrusted sources.</p>
eventname	VARCHAR(255) NULL	Specifies the name of the audit event. For example, AM-SESSION-CREATED and AM-SESSION-DESTROYED . For a list of audit event names, see "Audit Log Event Names".

Column	Datatype	Description
userid	VARCHAR(255) NULL	Specifies the universal identifier for authenticated users. For example, <code>id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com</code> .
trackingids	MEDIUMTEXT	Specifies the tracking IDs of the event, used by all topics.
runas	VARCHAR(255) NULL	Specifies the user to run the activity as. May be used in delegated administration. For example, <code>uid=amAdmin,ou=People,dc=example,dc=com</code> .
objectid	VARCHAR(255) NULL	Specifies the identifier of a system object that has been created, modified, or deleted. For example, <code>ou=SamuelTwo,ou=default,ou=OrganizationConfig,ou=1.0, ou=iPlanetAMAuthSAML2Service,ou=services,o=shop,ou=services, dc=example,dc=com</code> .
operation	VARCHAR(255) NULL	Specifies the state change operation invoked: CREATE , MODIFY , or DELETE .
beforeObject	MEDIUMTEXT NULL	Specifies the JSON representation of the object prior to the activity. For example: <pre> { "sunmspriority":["0"], "objectclass":["top", "sunServiceComponent", "organizationalUnit"], "ou":["SamuelTwo"], "sunserviceID":["serverconfig"] }</pre>
afterObject	MEDIUMTEXT NULL	Specifies the JSON representation of the object after the activity. For example:

Column	Datatype	Description
		<pre>{ "sunKeyValue": ["forgerock-am-auth-saml2-auth-level=0", "forgerock-am-auth-saml2-meta-alias=/sp", "forgerock-am-auth-saml2-entity-name=http://", "forgerock-am-auth-saml2-authn-context-decl-ref=", "forgerock-am-auth-saml2-force-authn=none", "forgerock-am-auth-saml2-is-passive=none", "forgerock-am-auth-saml2-login-chain=", "forgerock-am-auth-saml2-auth-comparison=none", "forgerock-am-auth-saml2-req-binding= urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect", "forgerock-am-auth-saml2-binding= urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact", "forgerock-am-auth-saml2-authn-context-class-ref=", "forgerock-am-auth-saml2-slo-relay=http://", "forgerock-am-auth-saml2-allow-create=false", "forgerock-am-auth-saml2-name-id-format= urn:oasis:names:tc:SAML:2.0:nameid-format:persistent"] }</pre>
changedfields	VARCHAR(255) NULL	Specifies the columns that were changed. For example, ["sunKeyValue"].
rev	VARCHAR(255) NULL	Not used.
component	VARCHAR(255) NULL	Specifies the AM service utilized. For example, Server Info , Users , Config , Session , Authentication , Policy , OAuth .
realm	VARCHAR(255) NULL	Specifies the realm where the operation occurred. For example, the Top Level Realm ("/") or the sub-realm name ("/shop").

am_auditconfig

am_auditconfig

Column	Datatype	Description
id	VARCHAR(56) NOT NULL	Specifies a universally unique identifier (UUID) for the message object, such as a568d4fe-d655-49a8-8290-bfc02095bec9-491 .
timestamp_	VARCHAR(29) NULL	Specifies the timestamp when AM logged the message, in UTC format to millisecond precision: yyyy-MM-ddTHH:mm:ss.msZ. For example: 2015-11-14T00:16:04.653Z
transactionid	VARCHAR(255) NULL	Specifies the UUID of the transaction, which identifies an external request when it comes into the system boundary. Any events generated while handling that request will be assigned that transaction ID, so that you may see the same transaction ID for different audit event topics. For example, 9c9e8d5c-2941-4e61-9c3c-8a990088e801 .

Column	Datatype	Description
		AM supports a feature where a trusted AM deployment with multiple instances, components, and ForgeRock products can propagate a transaction ID through each call across the stack. AM reads the X-ForgeRock-TransactionId HTTP header and appends an integer to the transaction ID. Note that this feature is disabled by default. When enabled, this feature should filter the X-ForgeRock-TransactionId HTTP header for connections from untrusted sources.
eventname	VARCHAR(255) NULL	Specifies the name of the audit event. For example, AM-CONFIG-CHANGE . For a list of audit event names, see "Audit Log Event Names".
userid	VARCHAR(255) NULL	Specifies the universal identifier for authenticated users. For example, id=scarter,ou=user,o=shop,ou=services,dc=example,dc=com .
trackingids	MEDIUMTEXT	Specifies the tracking IDs of the event, used by all topics.
runas	VARCHAR(255) NULL	Specifies the user to run the activity as. May be used in delegated administration. For example, uid=amAdmin,ou=People,dc=example,dc=com .
objectid	VARCHAR(255) NULL	Specifies the identifier of a system object that has been created, modified, or deleted. For example, ou=SamuelTwo,ou=default,ou=OrganizationConfig,ou=1.0, ou=iPlanetAMAuthSAML2Service,ou=services,o=shop,ou=services, dc=example,dc=com .
operation	VARCHAR(255) NULL	Specifies the state change operation invoked: CREATE , MODIFY , or DELETE .
beforeObject	MEDIUMTEXT NULL	Specifies the JSON representation of the object prior to the activity. For example: <div> <pre> { "sunsmpriority":["0"], "objectclass":["top", "sunServiceComponent", "organizationalUnit"], "ou":["SamuelTwo"], "sunserviceID":["serverconfig"] } </pre> </div>
afterObject	MEDIUMTEXT NULL	Specifies the JSON representation of the object after the activity. For example:

Column	Datatype	Description
		<pre>{ "sunKeyValue": ["forgerock-am-auth-saml2-auth-level=0", "forgerock-am-auth-saml2-meta-alias=/sp", "forgerock-am-auth-saml2-entity-name=http://", "forgerock-am-auth-saml2-authn-context-decl-ref=", "forgerock-am-auth-saml2-force-authn=none", "forgerock-am-auth-saml2-is-passive=none", "forgerock-am-auth-saml2-login-chain=", "forgerock-am-auth-saml2-auth-comparison=none", "forgerock-am-auth-saml2-req-binding= urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect", "forgerock-am-auth-saml2-binding= urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact", "forgerock-am-auth-saml2-authn-context-class-ref=", "forgerock-am-auth-saml2-slo-relay=http://", "forgerock-am-auth-saml2-allow-create=false", "forgerock-am-auth-saml2-name-id-format= urn:oasis:names:tc:SAML:2.0:nameid-format:persistent"] }</pre>
changedfields	VARCHAR(255) NULL	Specifies the columns that were changed. For example, ["sunKeyValue"].
rev	VARCHAR(255)	Not used.
component	VARCHAR(255) NULL	Specifies the AM service utilized. For example, Server Info , Users , Config , Session , Authentication , Policy , OAuth .
realm	VARCHAR(255) NULL	Specifies the realm where the operation occurred. For example, the Top Level Realm ("/") or the sub-realm name ("/shop").

Customizing CTS-Based Session Quota Exhaustion Actions

This section demonstrates a custom session quota exhaustion action plugin. AM calls a session quota exhaustion action plugin when a user tries to open more CTS-based sessions than their quota allows. Note that session quotas are not available for client-based sessions.

You only need a custom session quota exhaustion action plugin if the built-in actions are not flexible enough for your deployment. See ["Configuring Session Quotas"](#).

Creating & Installing a Custom Session Quota Exhaustion Action

You build custom session quota exhaustion actions into a .jar that you then plug in to AM. You must also add your new action to the Session service configuration, and restart AM in order to be able to configure it for your use.

Your custom session quota exhaustion action implements the [com.ipianet.dpro.session.service.QuotaExhaustionAction](#) interface, overriding the [action](#) method. The [action](#) method performs the action

when the session quota is met, and returns `true` only if the request for a new session should *not* be granted.

The example in this section simply removes the first session it finds as the session quota exhaustion action.

```
/*
 * The contents of this file are subject to the terms of the Common Development and
 * Distribution License (the License). You may not use this file except in compliance with the
 * License.
 *
 * You can obtain a copy of the License at legal/CDDLv1.0.txt. See the License for the
 * specific language governing permission and limitations under the License.
 *
 * When distributing Covered Software, include this CDDL Header Notice in each file and include
 * the License file at legal/CDDLv1.0.txt. If applicable, add the following below the CDDL
 * Header, with the fields enclosed by brackets [] replaced by your own identifying
 * information: "Portions copyright [year] [name of copyright owner]".
 *
 * Copyright 2012-2019 ForgeRock AS. All Rights Reserved
 */

package org.forgerock.openam.examples.quotaexhaustionaction;

import java.util.Map;

import javax.inject.Inject;

import org.forgerock.guice.core.InjectorHolder;
import org.forgerock.openam.session.Session;
import org.forgerock.openam.session.clientsdk.SessionCache;
import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

import com.iplanet.dpro.session.SessionException;
import com.iplanet.dpro.session.SessionID;
import com.iplanet.dpro.session.service.QuotaExhaustionAction;
import com.iplanet.dpro.session.service.SessionService;

/**
 * This is a sample {@link QuotaExhaustionAction} implementation,
 * which randomly kills the first session it finds.
 */
public class SampleQuotaExhaustionAction implements QuotaExhaustionAction {

    private static Logger debug = LoggerFactory.getLogger(SampleQuotaExhaustionAction.class);

    private final SessionCache sessionCache;
    private final SessionService sessionService;

    public SampleQuotaExhaustionAction() {
        this.sessionCache = InjectorHolder.getInstance(SessionCache.class);
        this.sessionService = InjectorHolder.getInstance(SessionService.class);
    }

    @Inject
    public SampleQuotaExhaustionAction(SessionCache sessionCache, SessionService sessionService) {
        this.sessionCache = sessionCache;
    }
}
```

```

        this.sessionService = sessionService;
    }

    /**
     * Check if the session quota for a given user has been exhausted and
     * if so perform the necessary actions. This implementation randomly
     * destroys the first session it finds.
     *
     * @param is The InternalSession to be activated.
     * @param existingSessions All existing sessions that belong to the same
     *                          uuid (Map:sid->expiration_time).
     * @return true If the session activation request should be rejected,
     *              otherwise false.
     */
    @Override
    public boolean action(
        Session is,
        Map<String, Long> existingSessions, long excessSessionCount) {
        for (Map.Entry<String, Long> entry : existingSessions.entrySet()) {
            try {
                // Get a Session from the cache based on the session ID, and destroy it.
                SessionID sessionId = new SessionID(entry.getKey());
                Session session = sessionCache.getSession(sessionId);
                sessionService.destroySession(session, sessionId);
                // Only destroy the first session.
                break;
            } catch (SessionException se) {
                if (debug.isDebugEnabled()) {
                    debug.debug("Failed to destroy existing session.", se);
                }
                // In this case, deny the session activation request.
                return true;
            }
        }
        return false;
    }
}

```

If you have not already done so, download and build the sample code.

For information on downloading and building AM sample source code, see [How do I access and build the sample code provided for AM \(All versions\)?](#) in the *Knowledge Base*.

In the sources, you find the following files:

[pom.xml](#)

Apache Maven project file for the module

This file specifies how to build the sample plugin, and also specifies its dependencies on AM components and on the Servlet API.

[src/main/java/org/forgerock/openam/examples/quotaexhaustionaction/SampleQuotaExhaustionAction.java](#)

Core class for the sample quota exhaustion action plugin

Once built, copy the .jar to `WEB-INF/lib/` where AM is deployed.

```
$ cp target/*.jar /path/to/tomcat/webapps/openam/WEB-INF/lib/
```

Using the `ssoadm` command, update the Session Service configuration:

```
$ ssoadm \
  set-attr-choicevals \
    --adminid uid=amAdmin,ou=People,dc=openam,dc=forgerock,dc=org \
    --password-file /tmp/pwd.txt \
    --servicename iPlanetAMSessionService \
    --schematype Global \
    --attributename iplanet-am-session-constraint-handler \
    --add \
    --choicevalues myKey=\
    org.forgerock.openam.examples.quotaexhaustionaction.SampleQuotaExhaustionAction
Choice Values were set.
```

Extract `amSession.properties` and if necessary the localized versions of this file from `openam-core-7.0.2.jar` to `WEB-INF/classes/` where AM is deployed. For example, if AM is deployed under `/path/to/tomcat/webapps/openam`, then you could run the following commands.

```
$ cd /path/to/tomcat/webapps/openam/WEB-INF/classes/
$ jar -xvf ../lib/openam-core-7.0.2.jar amSession.properties
inflated: amSession.properties
```

Add the following line to `amSession.properties`.

```
myKey=Randomly Destroy Session
```

Restart AM or the container in which it runs.

You can now use the new session quota exhaustion action. In the AM console, navigate to Configure > Global Services, click Session, scroll to Resulting behavior if session quota exhausted, and then choose an option.

Before moving to your test and production environments, be sure to add your .jar file and updates to `amSession.properties` into a custom .war file that you can then deploy. You must still update the Session service configuration in order to use your custom session quota exhaustion action.

Listing Session Quota Exhaustion Actions

List session quota exhaustion actions by using the `ssoadm` command:

```
$ ssoadm \
get-attr-choicevals \
--adminid uid=amAdmin,ou=People,dc=openam,dc=forgerock,dc=org \
--password-file /tmp/pwd.txt \
--servicename iPlanetAMSessionService \
--schematype Global \
--attributename iplanet-am-session-constraint-handler
```

I18n Key	Choice Value
choiceDestroyOldSession	org...session.service.DestroyOlddestAction
choiceDenyAccess	org...session.service.DenyAccessAction
choiceDestroyNextExpiring	org...session.service.DestroyNextExpiringAction
choiceDestroyAll	org...session.service.DestroyAllAction
myKey	org...examples...SampleQuotaExhaustionAction

Removing a Session Quota Exhaustion Action

Remove a session quota exhaustion action by using the **ssoadm** command:

```
$ ssoadm \
remove-attr-choicevals \
--adminid uid=amAdmin,ou=People,dc=openam,dc=forgerock,dc=org \
--password-file /tmp/pwd.txt \
--servicename iPlanetAMSessionService \
--schematype Global \
--attributename iplanet-am-session-constraint-handler \
--choicevalues \
org.forgerock.openam.examples.quotaexhaustionaction.SampleQuotaExhaustionAction
```

Choice Values were removed.

Glossary

Access control	Control to grant or to deny access to a resource.
Account lockout	The act of making an account temporarily or permanently inactive after successive authentication failures.
Actions	Defined as part of policies, these verbs indicate what authorized identities can do to resources.
Advice	In the context of a policy decision denying access, a hint to the policy enforcement point about remedial action to take that could result in a decision allowing access.
Agent administrator	User having privileges only to read and write agent profile configuration information, typically created to delegate agent profile creation to the user installing a web or Java agent.
Agent authenticator	Entity with read-only access to multiple agent profiles defined in the same realm; allows an agent to read web service profiles.
Application	<p>In general terms, a service exposing protected resources.</p> <p>In the context of AM policies, the application is a template that constrains the policies that govern access to protected resources. An application can have zero or more policies.</p>
Application type	<p>Application types act as templates for creating policy applications.</p> <p>Application types define a preset list of actions and functional logic, such as policy lookup and resource comparator logic.</p>

	Application types also define the internal normalization, indexing logic, and comparator logic for applications.
Attribute-based access control (ABAC)	Access control that is based on attributes of a user, such as how old a user is or whether the user is a paying customer.
Authentication	The act of confirming the identity of a principal.
Authentication chaining	A series of authentication modules configured together which a principal must negotiate as configured in order to authenticate successfully.
Authentication level	Positive integer associated with an authentication module, usually used to require success with more stringent authentication measures when requesting resources requiring special protection.
Authentication module	AM authentication unit that handles one way of obtaining and verifying credentials.
Authorization	The act of determining whether to grant or to deny a principal access to a resource.
Authorization Server	In OAuth 2.0, issues access tokens to the client after authenticating a resource owner and confirming that the owner authorizes the client to access the protected resource. AM can play this role in the OAuth 2.0 authorization framework.
Auto-federation	Arrangement to federate a principal's identity automatically based on a common attribute value shared across the principal's profiles at different providers.
Bulk federation	Batch job permanently federating user profiles between a service provider and an identity provider based on a list of matched user identifiers that exist on both providers.
Circle of trust	Group of providers, including at least one identity provider, who have agreed to trust each other to participate in a SAML v2.0 provider federation.
Client	In OAuth 2.0, requests protected web resources on behalf of the resource owner given the owner's authorization. AM can play this role in the OAuth 2.0 authorization framework.
Client-based OAuth 2.0 tokens	After a successful OAuth 2.0 grant flow, AM returns a token to the client. This differs from CTS-based OAuth 2.0 tokens, where AM returns a <i>reference</i> to token to the client.
Client-based sessions	AM sessions for which AM returns session state to the client after each request, and require it to be passed in with the subsequent

request. For browser-based clients, AM sets a cookie in the browser that contains the session information.

For browser-based clients, AM sets a cookie in the browser that contains the session state. When the browser transmits the cookie back to AM, AM decodes the session state from the cookie.

Conditions

Defined as part of policies, these determine the circumstances under which which a policy applies.

Environmental conditions reflect circumstances like the client IP address, time of day, how the subject authenticated, or the authentication level achieved.

Subject conditions reflect characteristics of the subject like whether the subject authenticated, the identity of the subject, or claims in the subject's JWT.

Configuration datastore

LDAP directory service holding AM configuration data.

Cross-domain single sign-on (CDSSO)

AM capability allowing single sign-on across different DNS domains.

CTS-based OAuth 2.0 tokens

After a successful OAuth 2.0 grant flow, AM returns a *reference* to the token to the client, rather than the token itself. This differs from [client-based OAuth 2.0 tokens](#), where AM returns the entire token to the client.

CTS-based sessions

AM [sessions](#) that reside in the Core Token Service's token store. CTS-based sessions might also be cached in memory on one or more AM servers. AM tracks these sessions in order to handle events like logout and timeout, to permit session constraints, and to notify applications involved in SSO when a session ends.

Delegation

Granting users administrative privileges with AM.

Entitlement

Decision that defines which resource names can and cannot be accessed for a given identity in the context of a particular application, which actions are allowed and which are denied, and any related advice and attributes.

Extended metadata

Federation configuration information specific to AM.

Extensible Access Control Markup Language (XACML)

Standard, XML-based access control policy language, including a processing model for making authorization decisions based on policies.

Federation

Standardized means for aggregating identities, sharing authentication and authorization data information between trusted providers, and

allowing principals to access services across different providers without authenticating repeatedly.

Fedlet	Service provider application capable of participating in a circle of trust and allowing federation without installing all of AM on the service provider side; AM lets you create Java Fedlets.
Hot swappable	Refers to configuration properties for which changes can take effect without restarting the container where AM runs.
Identity	Set of data that uniquely describes a person or a thing such as a device or an application.
Identity federation	Linking of a principal's identity across multiple providers.
Identity provider (IDP)	Entity that produces assertions about a principal (such as how and when a principal authenticated, or that the principal's profile has a specified attribute value).
Identity repository	Data store holding user profiles and group information; different identity repositories can be defined for different realms.
Java agent	Java web application installed in a web container that acts as a policy enforcement point, filtering requests to other applications in the container with policies based on application resource URLs.
Metadata	Federation configuration information for a provider.
Policy	Set of rules that define who is granted access to a protected resource when, how, and under what conditions.
Policy agent	Java, web, or custom agent that intercepts requests for resources, directs principals to AM for authentication, and enforces policy decisions from AM.
Policy Administration Point (PAP)	Entity that manages and stores policy definitions.
Policy Decision Point (PDP)	Entity that evaluates access rights and then issues authorization decisions.
Policy Enforcement Point (PEP)	Entity that intercepts a request for a resource and then enforces policy decisions from a PDP.
Policy Information Point (PIP)	Entity that provides extra information, such as user profile attributes that a PDP needs in order to make a decision.
Principal	Represents an entity that has been authenticated (such as a user, a device, or an application), and thus is distinguished from other entities.

	When a Subject successfully authenticates, AM associates the Subject with the Principal.
Privilege	In the context of delegated administration, a set of administrative tasks that can be performed by specified identities in a given realm.
Provider federation	Agreement among providers to participate in a circle of trust.
Realm	AM unit for organizing configuration and identity information. Realms can be used for example when different parts of an organization have different applications and identity stores, and when different organizations use the same AM deployment. Administrators can delegate realm administration. The administrator assigns administrative privileges to users, allowing them to perform administrative tasks within the realm.
Resource	Something a user can access over the network such as a web page. Defined as part of policies, these can include wildcards in order to match multiple actual resources.
Resource owner	In OAuth 2.0, entity who can authorize access to protected web resources, such as an end user.
Resource server	In OAuth 2.0, server hosting protected web resources, capable of handling access tokens to respond to requests for such resources.
Response attributes	Defined as part of policies, these allow AM to return additional information in the form of "attributes" with the response to a policy decision.
Role based access control (RBAC)	Access control that is based on whether a user has been granted a set of permissions (a role).
Security Assertion Markup Language (SAML)	Standard, XML-based language for exchanging authentication and authorization data between identity providers and service providers.
Service provider (SP)	Entity that consumes assertions about a principal (and provides a service that the principal is trying to access).
Authentication Session	The interval while the user or entity is authenticating to AM.
Session	The interval that starts after the user has authenticated and ends when the user logs out, or when their session is terminated. For browser-based clients, AM manages user sessions across one or more applications by setting a session cookie. See also CTS-based sessions and Client-based sessions.

Session high availability	Capability that lets any AM server in a clustered deployment access shared, persistent information about users' sessions from the CTS token store. The user does not need to log in again unless the entire deployment goes down.
Session token	Unique identifier issued by AM after successful authentication. For a CTS-based sessions , the session token is used to track a principal's session.
Single log out (SLO)	Capability allowing a principal to end a session once, thereby ending her session across multiple applications.
Single sign-on (SSO)	Capability allowing a principal to authenticate once and gain access to multiple applications without authenticating again.
Site	<p>Group of AM servers configured the same way, accessed through a load balancer layer. The load balancer handles failover to provide service-level availability.</p> <p>The load balancer can also be used to protect AM services.</p>
Standard metadata	Standard federation configuration information that you can share with other access management software.
Stateless Service	<p>Stateless services do not store any data locally to the service. When the service requires data to perform any action, it requests it from a data store. For example, a stateless authentication service stores session state for logged-in users in a database. This way, any server in the deployment can recover the session from the database and service requests for any user.</p> <p>All AM services are stateless unless otherwise specified. See also Client-based sessions and CTS-based sessions.</p>
Subject	<p>Entity that requests access to a resource</p> <p>When an identity successfully authenticates, AM associates the identity with the Principal that distinguishes it from other identities. An identity can be associated with multiple principals.</p>
Identity store	Data storage service holding principals' profiles; underlying storage can be an LDAP directory service or a custom IdRepo implementation.
Web Agent	Native library installed in a web server that acts as a policy enforcement point with policies based on web page URLs.